Calhoun: The NPS Institutional Archive

DSpace Repository

Theses and Dissertations | 1. Thesis and Dissertation Collection, all items

1997-09

Management of autonomous systems in the Navy's Automated Digital Network System (ADNS)

Sullivan, James A

Monterey, California. Naval Postgraduate School

http://hdl.handle.net/10945/8645

Downloaded from NPS Archive: Calhoun

http://www.nps.edu/library

Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

NPS ARCHIVE
1997.09
SULLIVAN, J. A.

DUDLEY KNOX LIBRARY
NAVAL POSTGRADUATE SCHOOL
MONTEREY CA 93943-5101

NAVAL POSTGRADUATE SCHOOL
Monterey, California

THESIS

MANAGEMENT OF AUTONOMOUS SYSTEMS IN THE NAVY'S AUTOMATED DIGITAL NETWORK SYSTEM (ADNS)

by

James A. Sullivan

September, 1997

Thesis Advisor: Rex A. Buddenberg
Associate Advisor: Suresh Sridhar

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	September 1997	Master's Thesis

4. TITLE AND SUBTITLE : Management of Autonomous Systems in the Navy's Automated Digital Network System (ADNS)

5. FUNDING NUMBERS

6. AUTHOR(S)
Sullivan, James A.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Postgraduate School
Monterey, CA 93943-5000

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT ***(maximum 200 words)***

In an effort to create a more efficient, interoperable communications environment for its ships at sea the Navy has developed the Automated Digital Network System. Because of its recent introduction into the fleet and the evolving nature of the program there has not yet been any high level operational guidance provided for communications planners and managers. The major contribution of this thesis is to describe key issues fundamental to successful mission accomplishment. Operating in a network-centric environment represents a conceptual departure from standard Navy at-sea communications methods. The changes in thinking necessitated by this departure are presented to highlight the need for a new approach to communications management. Analysis of program design and implementation yielded the framework for the outline of system requirements and the management considerations necessary for effective operational employment. Reviews of fundamental concepts underlying the system and program origins are provided as background material.

14. SUBJECT TERMS	15. NUMBER OF PAGES
Autonomous System, AS, Automated Digital Network System, ADNS	98
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFI- CATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

Approved for public release; distribution is unlimited

MANAGEMENT OF AUTONOMOUS SYSTEMS IN THE NAVY'S AUTOMATED DIGITAL NETWORK SYSTEM (ADNS)

James A. Sullivan
Lieutenant Commander, United States Navy
B.E., State University of New York Maritime College, 1984

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

from the

NAVAL POSTGRADUATE SCHOOL
September, 1997

NPS Archive
1997.09
Sullivan, J. A.

11615
S89016
c.1

DUDLEY KNOX LIBRARY
NAVAL POSTGRADUATE SCHOOL
MONTEREY CA 93943-5101

ABSTRACT

In an effort to create a more efficient, interoperable communications environment for its ships at sea the Navy has developed the Automated Digital Network System. Because of its recent introduction into the fleet and the evolving nature of the program there has not yet been any high level operational guidance provided for communications planners and managers. The major contribution of this thesis is to describe key issues fundamental to successful mission accomplishment. Operating in a network-centric environment represents a conceptual departure from standard Navy at-sea communications methods. The changes in thinking necessitated by this departure are presented to highlight the need for a new approach to communications management. Analysis of program design and implementation yielded the framework for the outline of system requirements and the management considerations necessary for effective operational employment. Reviews of fundamental concepts underlying the system and program origins are provided as background material.

TABLE OF CONTENTS

I. INTRODUCTION 1

II. ROLE OF ADNS IN THE NAVY'S COMMUNICATION INFRASTRUCTURE 5
A. PROGRAM GUIDANCE 5
1. Joint Doctrine 5
2. Navy Doctrine 5
3. JMCOMS 6
B. INTEROPERABILITY 6

III. WHY ADNS REQUIRES A NEW PERSPECTIVE ON COMMUNICATIONS MANAGEMENT 9
A. NETWORK TOPOLOGY 9
1. Interdependency 9
2. Backbone Network Redundancy 11
B. MULTICAST POTENTIAL 11
C. NEW APPLICATIONS 12

IV. FUNDAMENTAL REQUIREMENTS OF ADNS AUTONOMOUS SYSTEMS 13
A. DEFINING THE GENERIC AUTONOMOUS SYSTEM 13
B. BASIC CONSIDERATIONS FOR ESTABLISHING AUTONOMOUS SYSTEM BOUNDARIES 14
1. Common Mission 14
2. Contiguous Backbone 15
3. External Boundaries 15
4. Traffic Volume 16
5. Mission Requirements 19
6. Overall Perspective 19

V. CONSIDERATIONS FOR EFFECTIVE MANAGEMENT OF AUTONOMOUS SYSTEMS 21
A. BEYOND THE GENERIC AS – A SAMPLE SCENARIO 21
B. MANAGEMENT CONSIDERATIONS 23
1. Mission Communications Planning 23
a. Establish a Command Relationship 26
b. Anticipate Mission Changes 27
c. Command and Control Considerations 28
d. Establish a Casualty Response Plan 29
2. Casualty Conditions 29
a. Loss of a Single Radio 30
b. Loss of a Backbone Network 30
c. Loss of a Ship 32
3. Specific Considerations for Mission Changes or Casualty Conditions 32

a. Designate Critical Applications....32
b. OTCC Location....33
c. OSPF Adjustments....34
(1) Metrics....34
(2) Priority....34
(3) Hello Interval....34

VI. CONCLUSIONS AND RECOMMENDATIONS....35
A. AREAS FOR FUTURE RESEARCH....36

APPENDIX A. ADNS FUNDAMENTALS....37
A. INTRODUCTION....37
1. What is ADNS?....37
2. What is ADNS good for?....38
a. Mobile Platforms....39
b. Alternative to Wire/Fiber Transmission....40
3. Why invest in ADNS?....40
a. Quality of Service....41
b. Cost Effective Bandwidth....41
c. Leverages the Existing Internet....41
d. Flexibility....41
4. How does ADNS work?....42
5. ADNS Advantages....43
a. Removing Humans From the Loop....43
b. Load Sharing....43
c. Optimal Use of Bandwidth....44
d. Communications Agility....44
e. Transparency of Installation and Use....44
f. Logistics....45
g. Ease of Upgrade....45
h. Single Point for Communications Management....45
i. Ability to Transmit All Types of Data....45
6. ADNS Disadvantages....46
a. Cost of Installation....46
B. ADNS OPERATIONAL DESCRIPTION....46
1. Overview....46
2. Network Features....49
a. Routing Protocols....49
(1) OSPF....49
(2) BGP4....49
b. Logical Organization....50
3. Key Features/Functions....53
a. Priority....53
(1) Priority Tables....53
(2) Determining Message Priority....54
(3) Message Transmission....54

b. Load Balancing....55
c. Congestion Control....55
(1) Load Sharing....56
(a) Restrictions....57
(b) Implementation....57
(2) Source Quench....57
d. TCP Duplicate Packet Transmission Problems....58
(1) TCP Duplicate Packet Rejection....58
4. Integrated Network Management....59
a. Overview....59
(1) Local Control Center....60
(a) Network Manager....60
(b) Distributed Manager....61
(c) Communication Automation Manager....62
(2) Autonomous System Control Center....63
(3) Network Operations Center....64
b. Network Management Tools....64
(1) Network Management System Software....65
(2) Third Party Applications....65
C. HARDWARE....66
1. LAN....66
2. Router....67
3. CRIU....67
4. CAP....67
5. Cryptographic Device....67
6. Modem....67
7. Connectivity Media....67

APPENDIX B. APPLICABLE ROUTING PROTOCOL CONCEPTS....69
A. DEFINITIONS....69
B. INTERNAL ROUTING....70
1. OSPF....70
a. General....70
b. The Link State Database and Routing Table....71
c. Link State Advertisements....72
d. Routing Protocol Types....72
e. Establishing a Connection....75
(1) Discovering Neighbors and Verifying Two-way Communications....75
(2) Electing the Designated Router (DR) and Establishing Adjacency....75
f. Network Maintenance....77
g. Packet Routing....78
2. MOSPF....78
a. General....78

b. Characteristics of MOSPF 79
C. EXTERNAL ROUTING 80
1. BGP4 80
a. General 80
b. BGP4 Message Types 80
c. Operation 81
d. Routing Decision Process 81

LIST OF REFERENCES 83

INITIAL DISTRIBUTION LIST 85

LIST OF ACRONYMS

ABR	Area Border Router
ADNS	Automated Digital Network System
ARG	Amphibious Readiness Group
AS	Autonomous System
ASCC	Autonomous System Control Center
BDR	Backup Designated Router
BGP4	Border Gateway Protocol Version 4
C4I	Command, Control, Communications, Computers and Intelligence
C4IFTW	C4I for the Warrior
CA	Challenge Athena
CAM	Communications Automation Manager
CAP	Channel Access Protocol
COE	Common Operating Environment
COMMPLAN	Communications Plan
COTS	Commercial-off-the-Shelf
CRIU	Channel Access Protocol (CAP) to Router Interface Unit
CTP	Common Tactical Picture
CVBG	Aircraft Carrier Battle Group
DEFCON	Defense Condition
DII	Defense Information Infrastructure
DMR	Digital Modular Radio
DMS	Defense Messaging System
DNS	Domain Name Service
DoD	Department of Defense
DR	Designated Router
EGP	Exterior Gateway Protocol
EHF	Extremely High Frequency
EMCON	Emission Control
GPS	Global Positioning System
HF	High Frequency
IGP	Interior Gateway Protocol
INM	Integrated Network Management
IP	Internet Protocol
ITP	Integrated Terminal Program
IW	Information Warfare

JCS	Joint Chiefs of Staff
JMCOMS	Joint Maritime Communications System
JMCIS	Joint Maritime Command Information System
JTA	Joint Technical Architecture
JTF	Joint Task Force
LAN	Local Area Network
LCC	Local Control Center
LSA	Link State Advertisement
MOSPF	Multicast Open Shortest Path First
NBMA	Non-broadcast Multi-access
NCTAMS	Naval Computer and Telecommunications Area Master Station
NMS	Network Management System
NOC	Navy Operations Center
OSI	Open Systems Interconnection
OSPF	Open Shortest Path First
QoS	Quality of Service
RF	Radio Frequency
RFC	Request For Comments
SATCOM	Satellite Communications
SHF	Super High Frequency
SIPRNET	Secret IP Router Network
SNMP	Simple Network Management Protocol
TCP	Transmission Control Protocol
UHF	Ultra High Frequency
VTC	Video Teleconference
WAN	Wide Area Network

I. INTRODUCTION

Since the end of the Cold War a major thrust of the Department of Defense (DoD) and each of the armed services has been information dominance. Inter-service (joint) interoperability acts as a key enabler of that goal. However, reaching that goal has required making significant changes in the existing communications infrastructure of all of the armed services. Toward that end each of service has developed a program to support reaching that goal. For the Navy and Marine Corps that vision is in (Copernicus...Forward, 1995).

"Currently, each service uses components of the information spectrum to obtain data; however, due to insufficient or non-interoperable communications links, the data is still not transferred seamlessly" (Copernicus...Forward, 1995). Each of the services has multiple "stovepipe" communications systems, most of which cannot communicate amongst themselves, let alone communicate with the other services. The Navy has taken a giant leap towards seamless data transfer with the Automated Digital Network System (ADNS). ADNS provides the hardware necessary to integrate multiple independent systems into one common communications network. Creating such an environment has, in addition to many other improvements, made seamless interoperability achievable through the application of Internet concepts and standards.

Enterprise-wide networking is a new concept for a service used to its independent systems. This new architecture requires a conceptual shift in our way of handling communications, from a stovepipe to a network-centric framework. The ADNS program

provides the hardware and the network management tools to operate the system but, as is required of any new program, it also requires operational level guidance describing how it can best be employed.

As part of the Copernicus vision ADNS is capable of supporting all levels of warfighter including the Composite Warfare Commander and Joint Task Force Commander by employing networks that are flexible in size and number in order to support customized command and control (Copernicus...Forward, 1995). To achieve this goal in the face of such a markedly different operating environment creates a need for redefining the communications planning and execution processes as well as command and control relationships. Instead of being concerned only with the status of each communications circuit independently, commanders must now have a broadened network wide perspective with a view toward network optimization.

Existing documentation provides the technical details describing how ADNS operates and discusses the need for the Navy to provide employment guidance. This thesis addresses the lack of high level employment doctrine by providing the reader with guidance for managing ADNS networks. Chapter II provides the historical background that drove the development of ADNS. Based on an analysis of the ADNS program's operational characteristics Chapter III makes some comparisons between operations under ADNS and existing communications systems. Describing the conceptual differences in the operation of Navy communications systems as a result of employment of ADNS highlights the need for a new method of communications management.

Chapters IV and V fill the employment guidance gap by providing planners and at-sea communications managers with the essentials of mission planning and execution required for operations employing ADNS. Existing documentation outlines some proposed ADNS employment guidance. Chapter IV discusses the rationale that should be used by communications planners when considering these recommendations. Chapter IV also provides some additional considerations not addressed in the existing documentation. Chapter V proposes some mission planning guidance and casualty considerations for the operational managers. Chapter V also provides alternative suggestions and rationale for some of the management functions proposed in existing documentation. The appendices provide the reader with the baseline knowledge of ADNS operation and routing protocol concepts necessary to support the main body of the thesis.

Background information for this thesis was obtained from various draft documents provided by the design personnel at Naval Command Control and Ocean Surveillance Center (NRaD), San Diego. Hardware level training provided by Thung Tran and documentation by Roger Casey, both of NRaD, contributed a great deal to the authors understanding of the fundamentals of ADNS. Appendix A is an adaptation of a document written by the author, LT Brian Rehard, USN and LT Eric Andalis, USN.

II. ROLE OF ADNS IN THE NAVY'S COMMUNICATIONS INFRASTRUCTURE

A. PROGRAM GUIDANCE

1. Joint Doctrine

Prompted by the experience gained in Desert Storm, the 1990's has become a decade in which the U.S. military has been dedicated to restructuring its Command, Control, Communications, Computers and Intelligence (C4I) architecture. Recognizing that existing "stovepipe systems hinder operational flexibility in an environment of uncertainty" (C4I for the Warrior, 1993) has spurred Department of Defense (DoD) wide initiatives to fix the problem. C4I for the Warrior (C4IFTW) announced the Joint Chiefs of Staff's (JCS) vision of Joint Task Forces (JTFs) operating in a battle space that is fully integrated, interoperable and operates in a Common Operating Environment (COE) which permits effective coordination up, down and across chains of command (C4I for the Warrior, 1993). To achieve this vision one of the keys to success of any new C4I initiatives is interoperability.

2. Navy Doctrine

(Copernicus...Forward, 1995) provides the Navy's strategy, developed in response to the JCS vision. Updating the original Copernicus concept and incorporating the operational perspective of (Forward... from the Sea, 1994), (Copernicus...Forward, 1995) highlights the need for "rapid and reliable connectivity". By outlining four essential functions of C4I (Connectivity, Common Tactical Picture (CTP), Sensor to Shooter, and Information Warfare (IW)) the Navy has created a vision of the tactical

environment of the future. All four C4I functions are interrelated but connectivity is the key to implementing the other three. Achieving connectivity means that there is a bandwidth managed network of nodes through which information in any form (i.e. voice, video, data or imagery) can be passed. (Copernicus...Forward, 1995).

3. JMCOMS

"The Joint Maritime Communications Strategy (JMCOMS) implements the communication component of the Navy's Copernicus vision" (JMCOMS Master Plan, 1997). The three JMCOMS program elements; ADNS, Digital Modular Radio (DMR)/Slice and the Integrated Terminal Program (ITP) are designed to provide "high capacity, flexible communications under control of the warfighter" (JMCOMS Master Plan, 1997). JMCOMS attempts to create an environment in which RF media are shared among users so that bandwidth can be assigned on demand (JMCOMS Master Plan, 1997)

DMR/Slice will support communications in the 100KHz to 2GHz range while ITP covers the 2GHz and above range. ADNS provides the multiplexing capability that links existing stovepipe systems to create a radio-based wide area network (Radio-WAN). By combining and more efficiently employing the bandwidth in each individual system ADNS is able to improve information flow. (JMCOMS Master Plan, 1997)

B. INTEROPERABILITY

Compliance with the Defense Information Infrastructure (DII) COE means that a system must meet technical environment and program standards to ensure compatibility

with other systems. The COE also specifies the use of commercial off-the-shelf (COTS) products. To meet these requirements all ADNS unique components are implemented on COTS hardware. By using open and Military Standards (MilStd) protocols ADNS takes the necessary first step towards promoting interoperability with other services (see Figures 2.1 and 2.2). A fundamental architectural consideration in the design of ADNS was its ability to interact with the existing Internet. As a result systems capable of using the Internet can communicate with installations operating ADNS. (JMCOMS Master Plan, 1997)

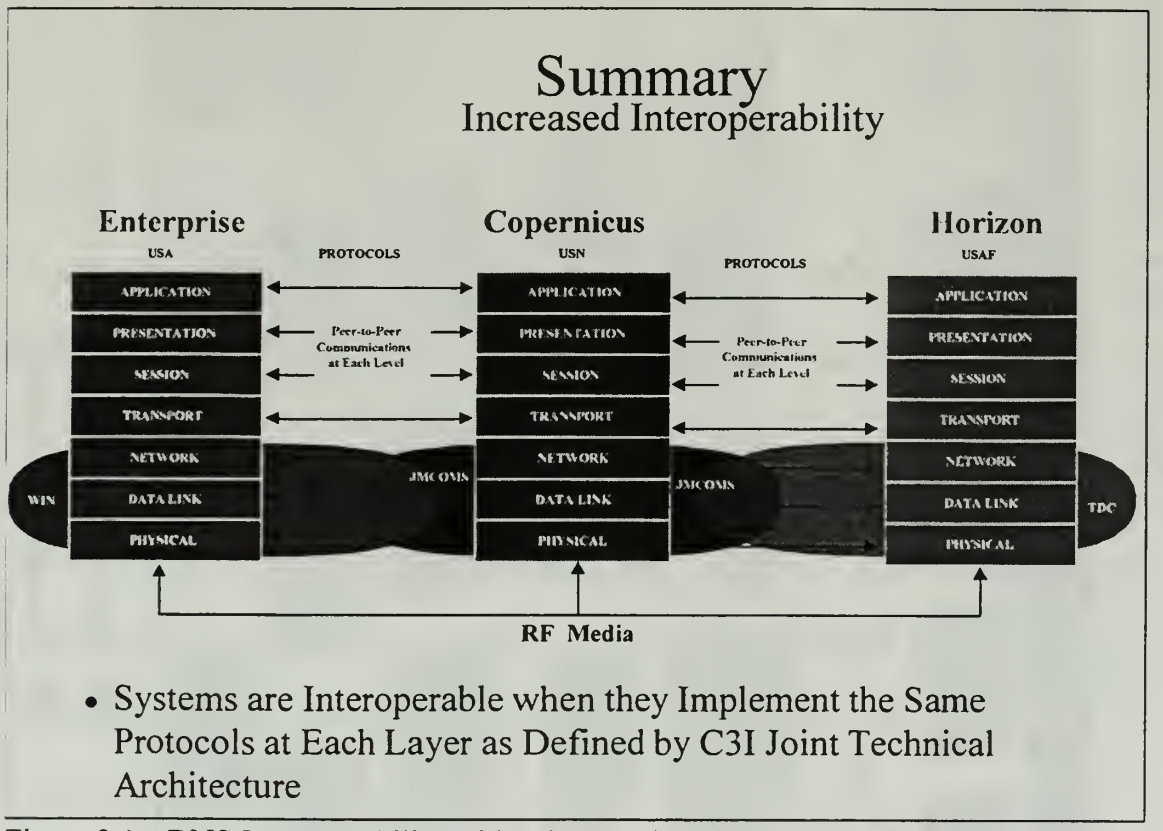

Figure 2.1 ADNS Interoperability with other services (From PIAC Brief, 1997)

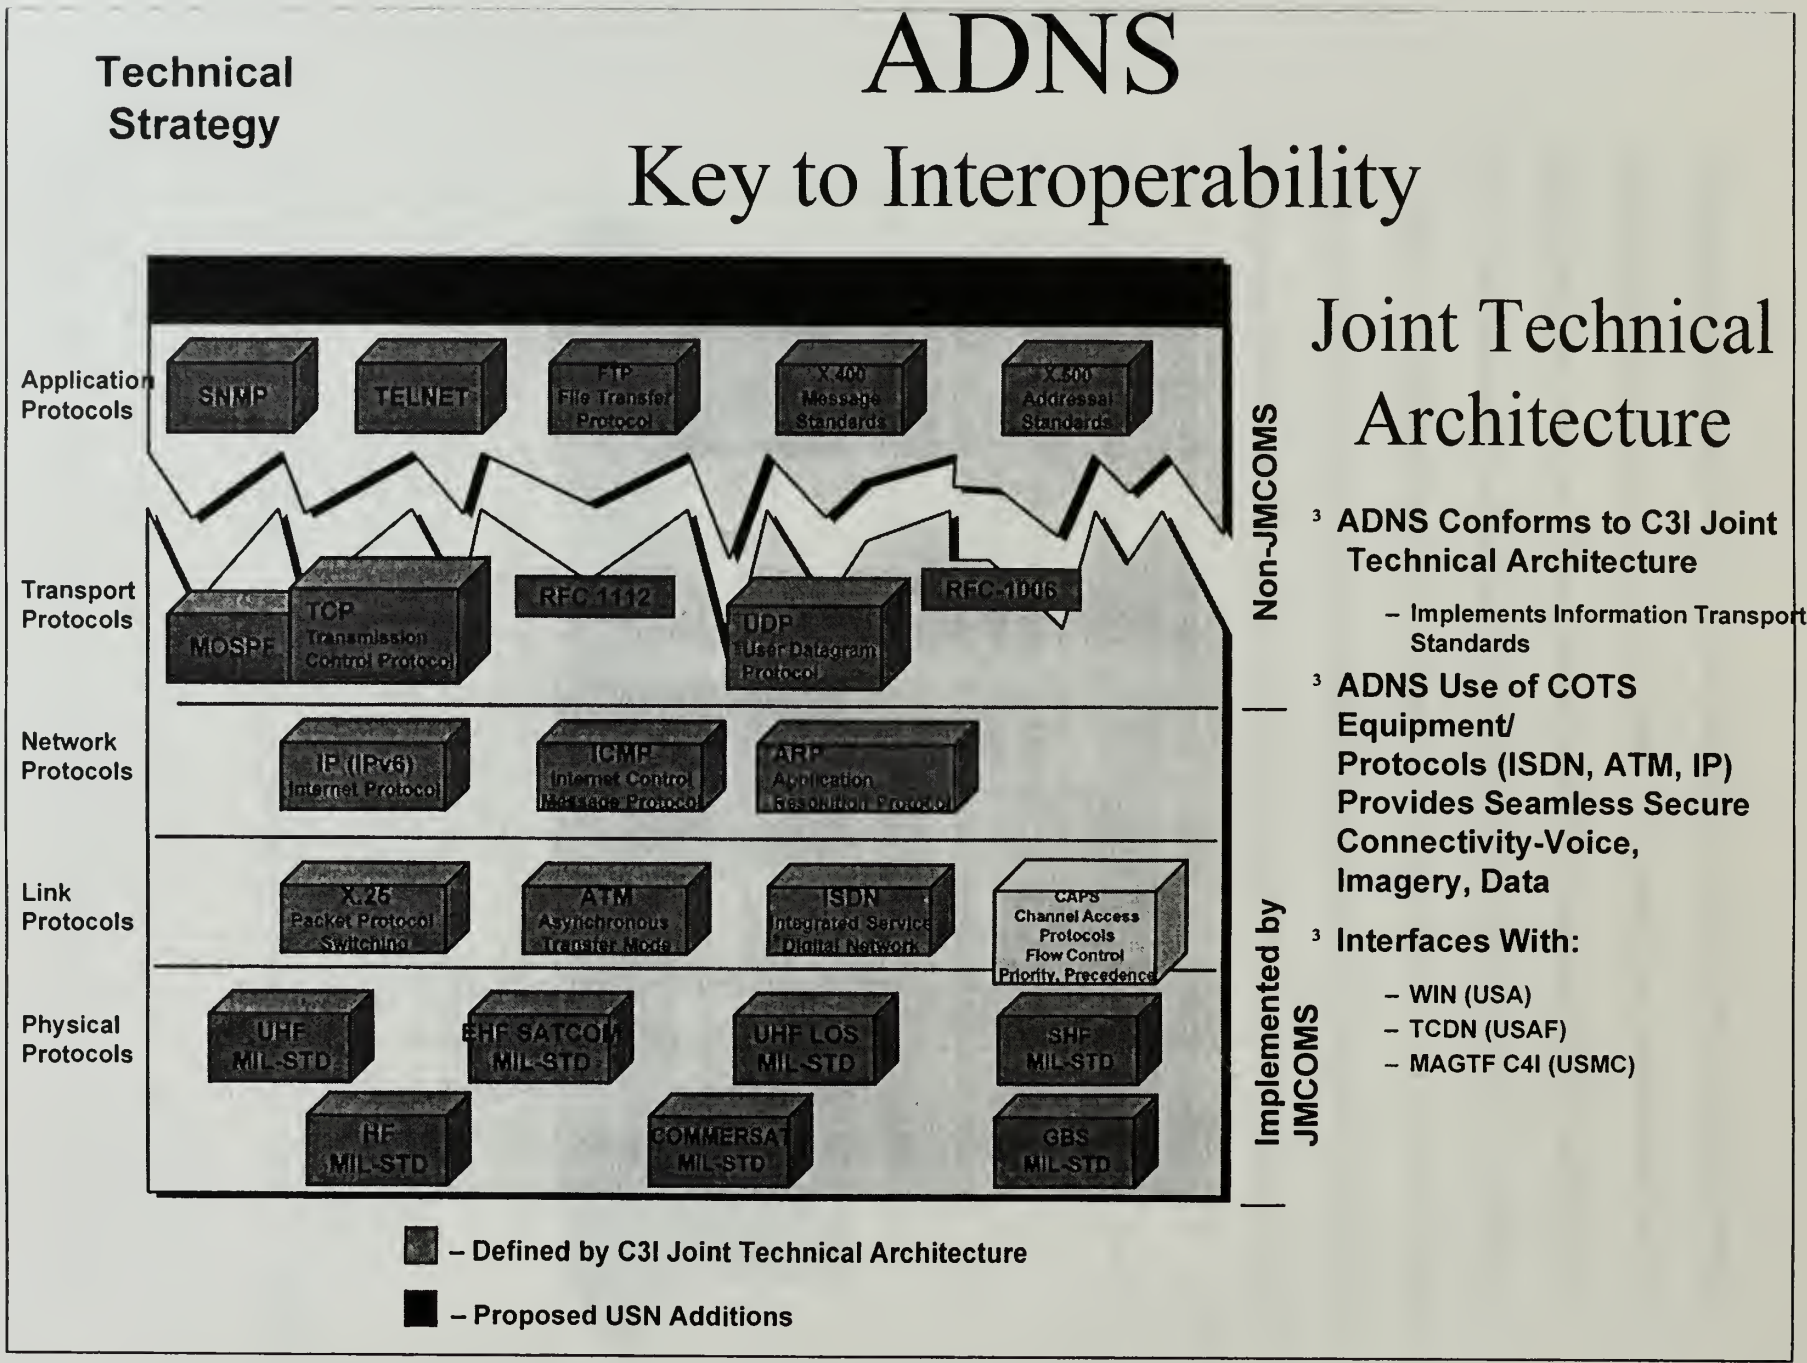

Figure 2.2 ADNS Protocols (From JMCOMS (SPAWAR PMW_176) Brief, 1997)

III. WHY ADNS REQUIRES A NEW PERSPECTIVE ON COMMUNICATIONS MANAGEMENT

A. NETWORK TOPOLOGY

Implementation of ADNS fundamentally changes the way Navy communications is accomplished. In a non-ADNS equipped ship, radio links are created to accomplish specific tasks. The task and the link are often inseparable. These task/link combinations were also for the most part independent of other task/link combinations. ADNS divorces the two, making the link transparent to the task.

When radio links are operated independently there is no need for network management because there is no network. ADNS creates a network and with it comes a need to change the historical methods of communications planning and circuit management.

1. Interdependency

An Autonomous System (AS) is " a group of routers exchanging routing information via a common routing protocol" (Moy, 1997). The ADNS AS is a network of different links interconnecting all of the attached installations. For ADNS to function efficiently there must be at least one link to each ship in the AS. Because the communications capabilities vary from ship to ship the level of participation in the network also varies. Smaller ships such as frigates or destroyers do not have as many communications circuits to dedicate to ADNS as larger platforms such as aircraft carriers

or cruisers (Casey, July 1997). As a result each ship is not connected on every circuit. This creates a situation where traffic may need to be routed through other units in order to reach its destination.

This third party relay function is important because it forces each platform to be aware of its own importance to the network. The network environment ADNS creates provides much greater redundancy and reliability for information transfer than previously existed with individual stovepipe systems. However with any system an effective implementation relies on understanding how the system operates. Previously if a radio link on a ship were to fail only that ship was affected. Now, with ADNS, one ship may be the sole relay station to one or more other ships in the AS. Loss of one circuit on a relaying ship could cause a complete loss of ADNS communications to several platforms. Proper planning should avoid this type of configuration if possible but each ship must still be aware of its relationship to the whole AS.

Additionally ships must be aware of this relay function because of its affect on Emission Control (EMCON). This is worthy of consideration for two reasons. First there is the potential for generating unwanted emissions. In a non-ADNS equipped platform transmissions are usually operator initiated or at least operator monitored and the source of the data is that ship. When acting as a relay platform transmissions can be initiated automatically in order to complete an exchange between two other members of the AS. A ship acting as an intermediate relay may be transmitting this exogenous traffic without operator intervention. Second, there is the potential for cutting off other ships

when a relay platform enters an EMCON condition. A relay platform could unknowingly isolate other units by terminating emissions on necessary relay links.

2. Backbone Network Redundancy

The combination of multiple links using different radio systems and the virtual link relay capability provided by ADNS forms a robust network with the potential for providing multiple redundant information flow paths. To optimize the flow of information this redundancy should be exploited wherever possible. To do so requires a detailed knowledge of not only the capabilities of the available radio systems but also real time knowledge of how those systems are configured with respect to ADNS.

B. MULTICAST POTENTIAL

ADNS provides the ability to multicast traffic within an AS. Although similar to a RF broadcast situation multicast on limited bandwidth links brings with it some important considerations. Among those considerations are the number of members being addressed and whether each unit is being reached directly or via third party relay. Because RF links do not operate at the same capacity as typical landline wire/fiber connections the use of multicast must be a constantly evaluated alternative. The available bandwidth is too limited to waste on inefficient practices. Multicast with too few members may not be as efficient as unicast. Also when relaying through third parties multicast may become less efficient than unicast.

C. NEW APPLICATIONS

ADNS deals strictly with IP datagrams. This means the range of applications that can be used in an ADNS network includes anything that can be transmitted across a standard IP network. In addition to military applications such as the Defense Messaging System (DMS) and Joint Maritime Command Information System (JMCIS), ADNS was specifically designed to support such functions as e-mail, file transfer and video teleconferencing. The ability to use these types of applications at sea is new to the Navy. What is also new is the access to these applications throughout the chain of command. Shipboard LANs are connected to ADNS, providing essentially any PC user with external communications capability. The impact of, for example, providing e-mail access to every member of a ship's crew needs to be evaluated as a part of our continuously evolving command and control architecture. ADNS provides a level of access never before experienced. Exploiting this access may, in some cases, be desirable. In other situations, it may be necessary to reconstruct, via policy implemented in hardware/software configuration, the barriers that ADNS has so effectively lowered.

IV. FUNDAMENTAL REQUIREMENTS OF ADNS AUTONOMOUS SYSTEMS

A. DEFINING THE GENERIC AUTONOMOUS SYSTEM

The Autonomous System is a routing protocol concept used to establish logical routing boundaries (Moy, 1997). For a review of routing protocol concepts applicable to ADNS see Appendix B. As a primarily afloat force, it is logical that a Navy AS is made up of ships. Each ship will usually, for internal routing purposes, be considered an area. An AS will typically have multiple ships and one or more shore stations as its members. The shore station, a Naval Communications and Telecommunications Area Master Station (NCTAMS), will act as the boundary for passing traffic to and from the AS. Since the Exterior Routing Protocol (ERP), BGP4, requires a stable environment the NCTAMS was chosen as its host (Casey, July 1997).

A generic AS can be viewed as "a collection of ships with one or more shore entry/exit points" as shown in Figure 4.1 (Casey, July 1997). In general the AS should consist of ships with a common mission and thus a need for routine, rapid ship to ship communications. (Casey, July 1997). The AS grouping is a logical vice geographical one.

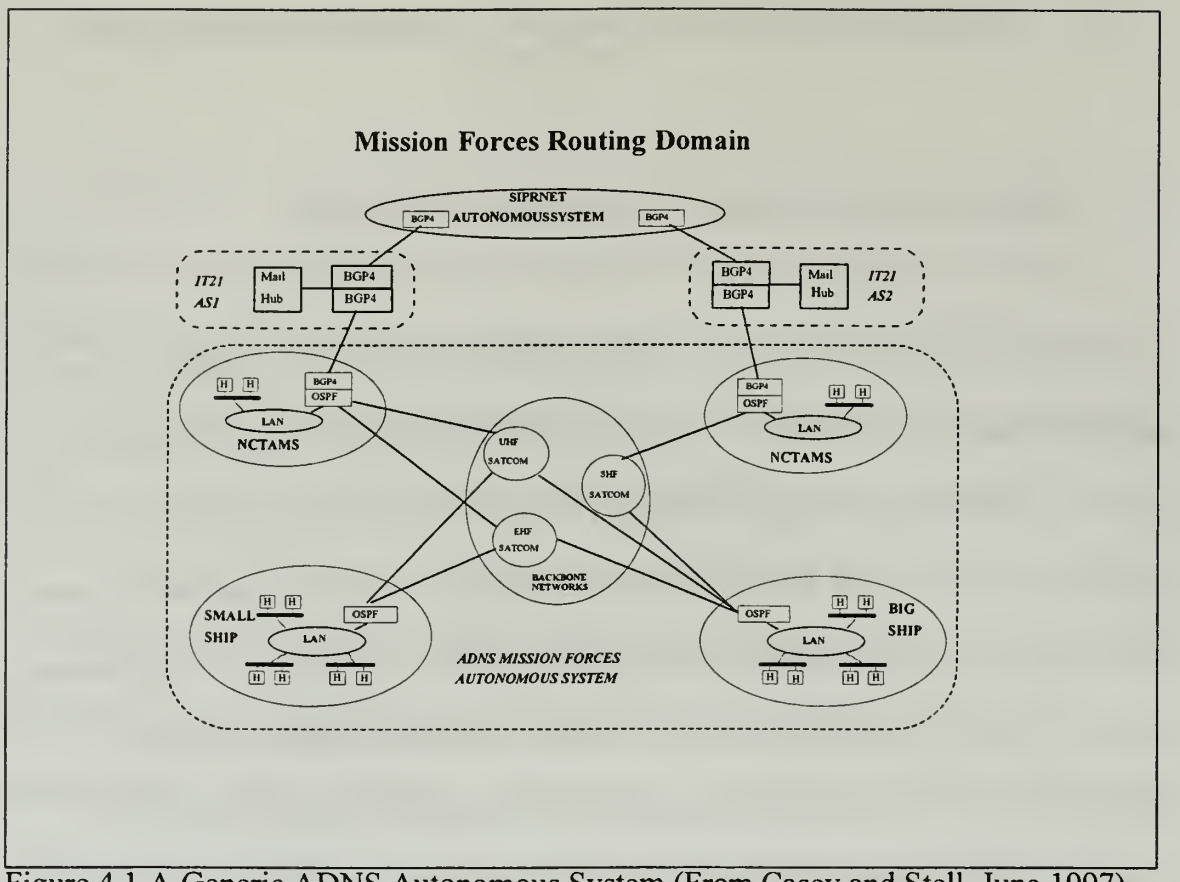

Figure 4.1 A Generic ADNS Autonomous System (From Casey and Stell, June 1997)

B. BASIC CONSIDERATIONS FOR ESTABLISHING AUTONOMOUS SYSTEM BOUNDARIES

1. Common Mission

One of the primary considerations when establishing an AS is that members of the network should share a common mission. The common mission means there is a need to share common data. It is a more efficient use of bandwidth if the source and all of the destination addressees are in one AS and none in an adjoining AS. It is also more efficient if every member of the AS is an Addressee on a message. Although it is important, it is only one of the factors that needs to be considered when establishing an AS.

2. Contiguous Backbone

Each platform that is to participate in the AS must have at least one network on which it can connect to the backbone. Although this may be an obvious requirement it still needs to be part of the planning process. The ships (and their respective radio suites) that are potential candidates for inclusion in an AS will determine the minimum number of platforms needed to form a contiguous network. Due to the limited selection of radios available on smaller platforms it may be necessary to establish virtual links to include all ships in the network.

3. External Boundaries

An AS will usually have one NCTAMS as its dedicated primary communications support facility. This is normally a geographic consideration. Each NCTAMS is responsible for a different part of the world. However this does not mean that a ship is not using the services of more than one NCTAMS at any given time. Each NCTAMS can participate in the broadcast circuits (such as UHF SATCOM and EHF SATCOM) available to many platforms. However several point-to-point communications systems such as Challenge Athena (CA) and SHF SATCOM are available only between a ship and a NCTAMS. This impacts the location of potential external boundaries because not all these links are available at every NCTAMS. Consequently the radio systems available on each platform will play a large part in deciding where and how many external boundaries there are in a given AS.

In addition to providing communications links the NCTAMS also perform other functions such as providing Domain Name Service (DNS) and acting as the mail hub for

the AS. Although these functions can be transferred, they can only be performed by a single NCTAMS at any given time

Since there can only be one advertised route into and out of an AS then a decision must be made regarding which of the connected shore stations will perform that function. Proper selection of the external gateway can greatly reduce the overhead incurred by platforms in the AS. (Casey and Stell, June 1997)

4. Traffic Volume

There are two factors to consider when discussing traffic volume. First is the contribution of overhead to total traffic volume. Overhead is the traffic being passed among the network's routers. This traffic is necessary and of high priority. Without accurate knowledge of the status of the network the routers cannot effectively route traffic.

The inter-router communications required to maintain the routing tables creates a substantial amount of traffic on the backbone network. The biggest contribution to overhead is from the Hello packets that are sent and received by all ABRs connected to the backbone (Casey and Stell, June 1997).

The impact of this overhead depends on the number of node on the network. Examination of Figure 4.2 shows that, for a given number of nodes on a network, the percentage of the total link capacity consumed by overhead varies with link capacity. Overhead uses a smaller percentage of the total capacity as link capacity increases. This means that for a given maximum overhead percentage a higher capacity link can support more nodes than a lower capacity link. (Casey and Stell, June 1997)

The significant part of this discussion is that the driving factor dictating the maximum number of nodes is not the highest capacity available to the AS. Instead the capability of each ship must be evaluated and compared to all others in the AS. The limiting platform is that ship whose highest capacity link is less than the highest capacity link of every other ship in the AS. For example suppose there are five ships being placed in an AS and four of them are capable of 64kbps but the fifth one is only capable of 2.4kbps. The limiting platform is the one capable of only 2.4kbps. Since all ABRs in the AS are passing data regardless of capacity the lower bandwidth link must still handle the traffic from all other ABRs. Thus it is the maximum capacity of the limiting platform that will limit the number of ships in the AS.

The other component of total traffic volume is data volume. Data includes all the packets being transmitted in support of any end users attached to the network. This consideration will also tend to drive the upper limit to the number of ships in an AS. The anticipated traffic volume should be considered with respect to the available capacity. If the anticipated volume will cause potential congestion problems it may be necessary to supersede the common mission consideration and form more than one AS. This may be a decision driven by the mix of radios available on each platform. Ships with limited bandwidth capability may need to be segregated to allow the higher bandwidth capable platforms to operate closer to capacity.

One negative aspect of splitting platforms with a common mission into more than one AS is the increase in traffic through the ASBRs at the NCTAMS. Depending on the volume of data passing between the ASs this may cause a loading problem for the ASBRs. Another effect of splitting into two systems is the duplication of information

Nodes in network 10

				Avg. % Bandwidth (BPS)		
	bits	Timers (sec)	avg. bps	2400	16000	32000
Router links	13,440	1800	7.47	0.31%	0.05%	0.02%
Network links	704	1800	0.39	0.02%	0.00%	0.00%
Summary links	7,360	1800	4.09	0.17%	0.03%	0.01%
Link State ACK	6,400	1800	3.56	0.15%	0.02%	0.01%
Total	27,904		15.50	0.65%	0.10%	0.05%
Peak time in seconds				11.63	1.74	0.87
Hello Packets	6,400	30	213.33	8.89%	1.33%	0.67%

Nodes in network 30

				Avg. % Bandwidth (BPS)		
	bits	Timers (sec)	avg. bps	2400	16000	32000
Router links	97,920	1800	54.40	2.27%	0.34%	0.17%
Network links	1,344	1800	0.75	0.03%	0.00%	0.00%
Summary links	41,280	1800	22.93	0.96%	0.14%	0.07%
Link State ACK	19,200	1800	10.67	0.44%	0.07%	0.03%
Total	159,744		88.75	3.70%	0.55%	0.28%
Peak time in seconds				66.56	9.98	4.99
Hello Packets	38,400	30	1,280.00	53.33%	8.00%	4.00%

Nodes in network 50

				Avg. % Bandwidth (BPS)		
	bits	Timers (sec)	avg. bps	2400	16000	32000
Router links	259,200	1800	144.00	6.00%	0.90%	0.45%
Network links	1,984	1800	1.10	0.05%	0.01%	0.00%
Summary links	100,800	1800	56.00	2.33%	0.35%	0.18%
Link State ACK	32,000	1800	17.78	0.74%	0.11%	0.06%
Total	393,984		218.88	9.12%	1.37%	0.68%
Peak time in seconds				164.16	24.62	12.31
Hello Packets	96,000	30	3,200.00	133.33%	20.00%	10.00%

Figure 4.2 Network OSPF Loading (From Casey and Stell, June 1997)

that may be necessary when sending identical traffic to platforms in both ASs. This situation results in the same type of wasted bandwidth problem that exists with current stovepipe systems, a problem ADNS was designed to avoid.

5. Mission Requirements

As discussed above, the mission will to a large degree drive membership in the AS. The consideration of mission related traffic volume should be considered along with the impact of the mission itself. During peacetime operations the potential for loss of members of the network is minimal. However when faced with hostilities there is the possibility of loss of individual radios, backbone subnets or even entire platforms. Consequently the ability to continue operations despite communications casualties is a necessity. To deal with this alternative the amount of redundancy within the backbone should be evaluated. Such things as single points of failure that will disrupt the contiguous nature of the backbone must be identified and contingency plans developed.

One alternative that may solve several problems is to consider including a platform in the AS that does not have a common mission. A ship with an extensive communications capability can have provide additional bandwidth as well as provide redundancy in the backbone. As long as the new ship does not bring with it an overwhelming communications requirement that will negatively impact traffic flow in the AS this is a reasonable alternative.

6. Overall Perspective

These five areas should be used as a general guide. Each situation will be different and the available alternatives must be evaluated based on the current conditions. As operational experience is gained with ADNS additional factors may be seen to play an

important role in this decision process. The important point is to develop and apply a network-centric view of this new environment in order to anticipate the demands it will make on the AS.

V. CONSIDERATIONS FOR EFECTIVE MANAGEMENT OF AUTONOMOUS SYSTEMS

As is true of any operation a certain amount of preplanning is required to increase its likelihood of success and to help ensure a smoother operation. This section attempts to point out some of the important considerations that should be used in the planning process to allow mission forces to better react to both planned events and casualty situations encountered during mission execution. The focus is on those areas that are of concern to and can be influenced by the operating forces within an AS. Specific actions are not provided since the appropriate action will be dictated by the specifics of a given event.

A. BEYOND THE GENERIC AS – A SAMPLE SCENARIO

The generic AS consists of multiple ships at sea with a common mission (Casey and Stell, June 1997). The generic AS is an adequate model to apply to a group of ships conducting an open ocean transit. An added level of complexity is encountered when a shift in missions occurs or divergent mission requirements make splitting into multiple ASs a viable option. For example a transiting task force could consist of a carrier battle group (CVBG) and an Amphibious Ready Group (ARG). While enroute to a destination they share a common mission. The need to communicate amongst platforms in order to share tactical information or weather etc. makes the decision to link these ships together in one AS a logical one.

However, once the destination is reached the ARG will likely break away and begin its task of conducting an amphibious landing. The CVBG, although operating in support of the same operation, has a very different mission (such as providing air support and/or naval gunfire support). The communications requirements of both groups have shifted with the mission shift. Although both groups will have the need to communicate with each other (inter-AS) the overwhelming portion of the communication will likely be amongst themselves (intra-AS). In fact the ARG's communications requirements will actually expand as its mission begins. The landing craft, such as LCACs and helicopters will establish and maintain communications with their host ships throughout the operation. In addition to voice communications requirements, these remote platforms like the LCAC have the potential to provide valuable tactical data, in the form of radar information, back to their host ship, or other larger ships in the ARG, standing well off shore.

Yet another logical shift could occur once the Marine landing force is established ashore. During the landing phase the Marine force is supported by the Navy landing force and their communications requirements could logically be grouped in that AS. However, when the Marine commander has shifted ashore there is less commonality in mission and having the Marine force establish its own AS is a logical extension of the intent to logically group by common mission requirements.

Doctrine hasn't been written yet regarding the adoption of ADNS by other services and its use in scenarios like this one. However, the shift to a network communications environment, such as that provided by ADNS, could significantly

improve mission effectiveness. The same advantages, such as redundancy and reliability, afforded the Navy by its shift to ADNS are available to the other services as well.

B. MANAGEMENT CONSIDERATIONS

An operation can be logically divided into two phases: mission planning and mission execution. The execution phase could again be subdivided into planned and unplanned (or casualty) events. If the mission planning process has been performed adequately the mission execution phase should simply be a matter of implementing the plan. Recognizing that missions can and do change, sometimes with little or no notice, it will not be possible to anticipate and plan for every contingency. When unforeseen situations arise, adapting on the fly may be necessary. In such situations it is important to first make sure communications are maintained and then at the earliest opportunity evaluate the situation with respect to the planning guidance to fill any gaps or correct any deficiencies that may exist.

1. Mission Communications Planning

Prior to any operation there are a number of issues that should be addressed and operating rules established. The result of this process is the Communications Plan (COMMPLAN). The operations planning process is outlined in Figure 5.1. The COMMPLAN encompasses both policy and hardware/software configuration issues. (Casey, July 1997)

For the purposes of this discussion, policy issues are those areas where the assignment of configuration parameters requires some decision process affecting network

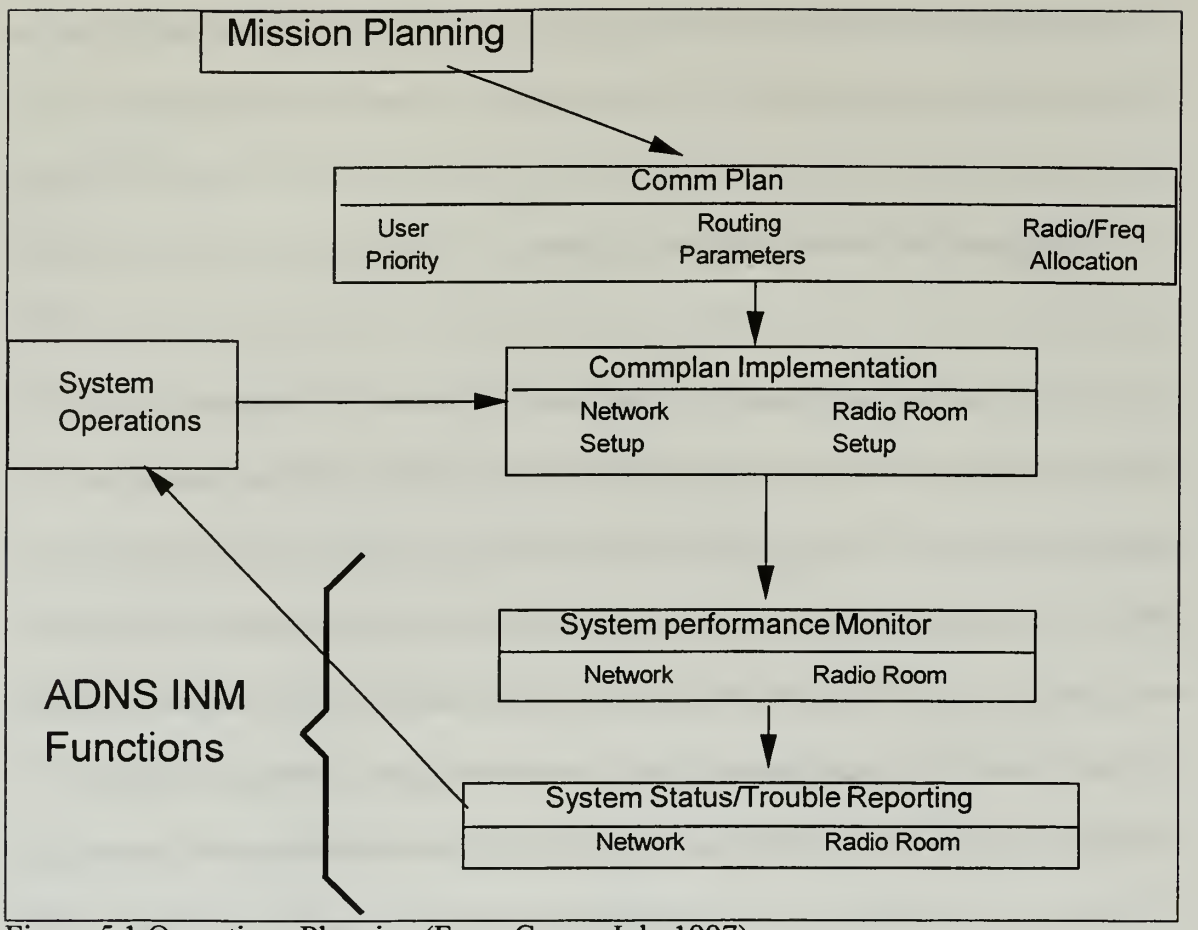

Figure 5.1 Operations Planning (From Casey, July 1997)

optimization, either within an AS or among multiple ASs. For example, the assignment of radio frequencies, while necessary, is not a policy decision. It does not generally impact communications flow beyond the requirement that everyone must know and use a designated frequency and that frequency does not suffer from or cause interference problems. The results of policy decisions may then implemented as hardware or software settings as required. Figures 5.2 and 5.3 show some of the areas addressed in the different phases of the planning and implementation processes. The following areas should be used as part f the mission planning/COMMPLAN formulation process.

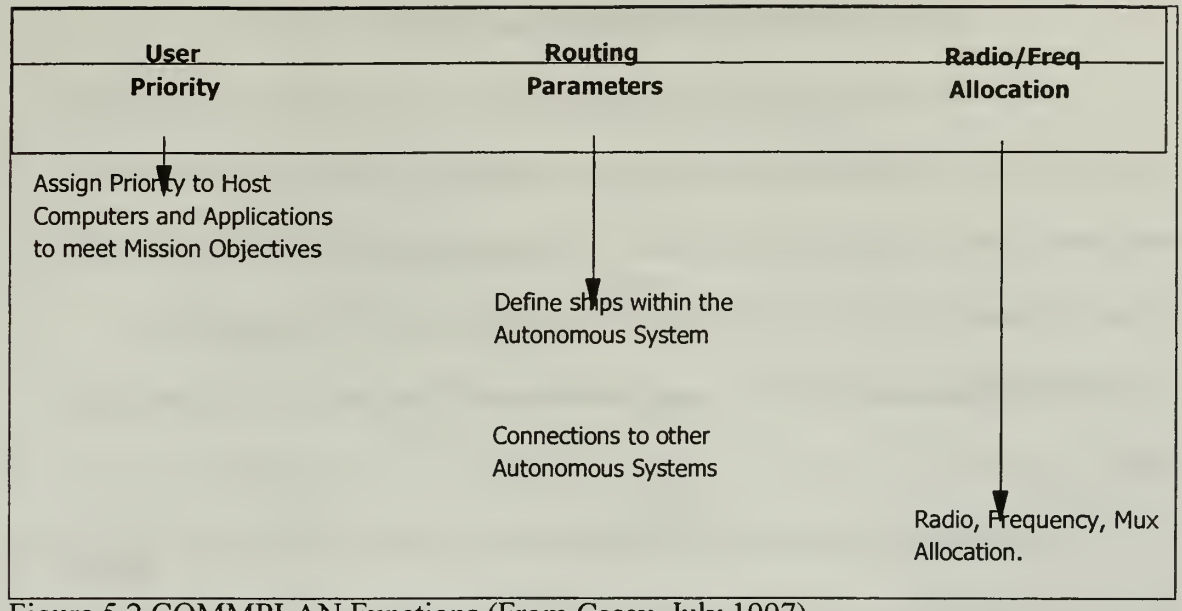
User Priority
Routing Parameters
Radio/Freq Allocation
Assign Priority to Host Computers and Applications to meet Mission Objectives
Define ships within the Autonomous System
Connections to other Autonomous Systems
Radio, Frequency, Mux Allocation.
Figure 5.2 COMMPLAN Functions (From Casey, July 1997)

Commplan Implementation

Network Setup	Radio Room Setup
Determine Routing Metric Values for OSPF Routing	Define Radios and Frequencies
Determine Policy for BGP4 Routing	Crypto Keys
Assign Priorities to user/ Applications (1 thru 15).	Set Mux allocations to IP networks and switched systems.
Set Queue Thresholds	Define Satellite resources and acquire channels.
Set IP addressing plan for exterior comms.	
Define Mail Hubs and DNS Requirements	

Figure 5.3 COMMPLAN Implementation (From Casey, July 1997)

a. Establish a Command Relationship

The ADNS program provides for a LCC on most ships and an ASCC at each NCTAMS. In addition to hardware configuration functions The LCC can also monitor the status of the backbone subnets through a graphical display similar to that shown in Figure 5.4. The ASCC will provide network monitoring functions for connections between multiple ASs via summary reports form the LCCs. (Casey, July 1997)

The focus of the LCC is on the individual ADNS installation. There is no provision for providing communications guidance for a single AS from within that AS. There should be a designated sea-based "Officer in Tactical Communications Command" (OTCC) that makes key communications decisions for the AS. He needs to be sea based to have access to the information necessary to maintain an understanding of the tactical situation. The NCTAMS is an adequate facility for communications management but its mission is not tactical.

(Casey, July 1997) also describes the higher level management functions to be performed by the NOC. Although there are several of those functions that impact multiple ASs (such as reassigning DAMA channels or reallocating bandwidth) there are also several functions that are better handled by a decision maker within the AS. Such parameters as metric value and priority assignment are AS specific issues that may need to be adjusted based on the tactical picture, which is not available at the NOC. (Casey, July 1997) in fact points out that the assignment of metric values to a given subnet could be different for different ships, depending on the role of that ship. When roles shift

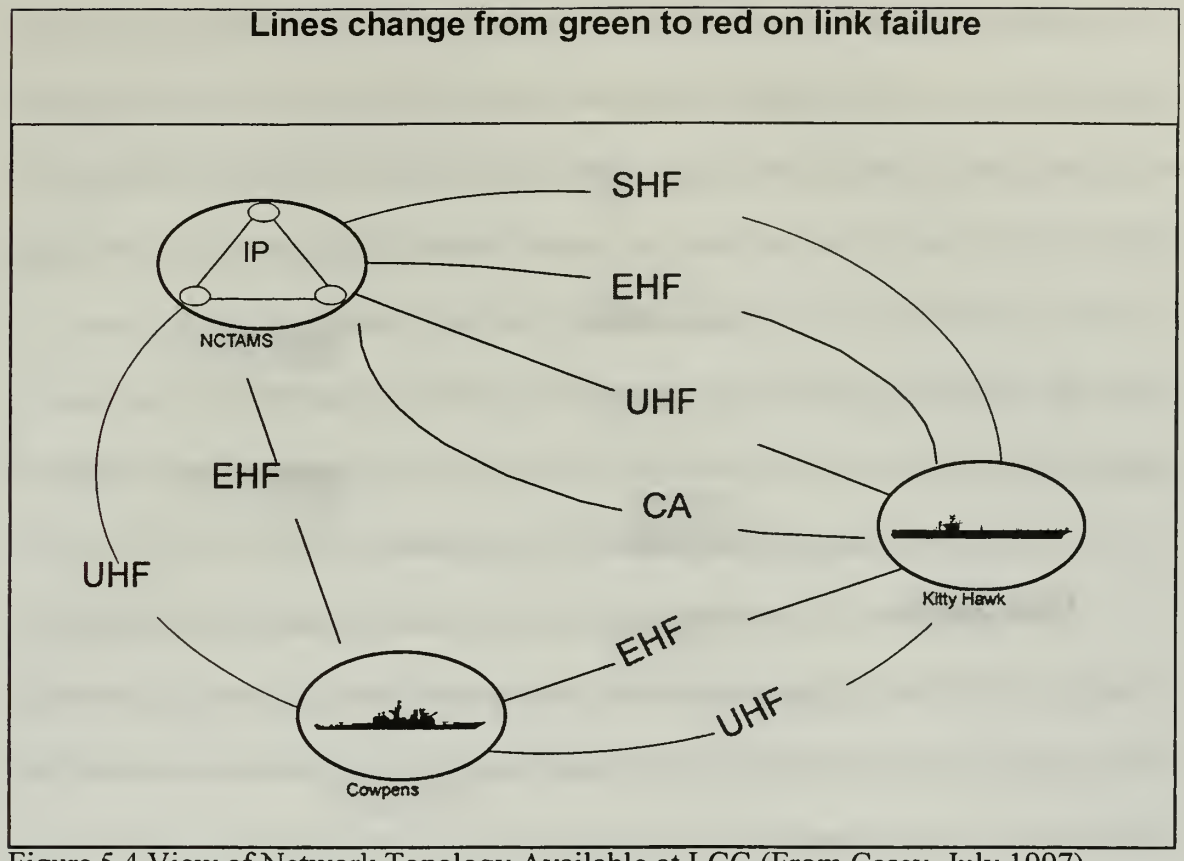

Figure 5.4 View of Network Topology Available at LCC (From Casey, July 1997)

within an AS the metrics and priorities may also need to shift. The inter-AS routing is essentially transparent to those outside the AS so a sensible configuration would be to control these functions from within. Internal control also reduces the traffic destined for outside the AS.

b. Anticipate Mission Changes

Arguably the most important aspect of making changes during a mission or operation is the amount of planning that has gone into making the plan itself. In general the better the planning process the better the transition. Almost as important is when to conduct the change. Regardless of the level of planning, completing an AS reorganization will cause a certain amount of disruption to the units participating in the change. For example, as discussed in the scenario above, when the ARG branches off

into its own AS there will be an increase in the amount of network administrative traffic generated as a result. Depending on their final location, electing a new DR and BDR will have to take place in one AS or the other. This process and the subsequent exchange of database information that occurs during the process of establishing adjacencies will cause an increase in overhead that may create unacceptable delays in the passing of mission critical data. Therefore a shift in AS organization should occur in advance of the actual mission shift to allow the network to stabilize before commencing with the new mission.

c. Command and Control Considerations

Because of the broad level of access to conduct external communications provided by ADNS each command must now develop policy dictating access rules. The ability to have any level in the chain of command communicate with other commands has both positive and negative aspects.

Effective implementation of the Sensor-to-Shooter concept may require lowering the command and control barriers currently in place. To conduct time sensitive missions it may be desirable to allow a more direct path between the information source and the weapons delivery system operator. Simultaneous reception of targeting information at all levels of the chain of command could minimize the processing time by instilling a command by negation policy. (Copernicus...Forward, 1995)

On the negative end of the spectrum, greater access can also result in low priority traffic generated at lower levels in the chain of command causing communications bottlenecks. Because of the network nature of ADNS these disruptions can impact more than one ship in the AS.

The question that must be addressed is what is the appropriate level of access or command and control for the specific task? ADNS has lowered the traditional access barriers. It is now a matter of generating the policies that raise those barriers to the appropriate level in certain areas.

d. Establish a Casualty Response Plan

ADNS creates a different communications environment than we are used to. Because of the sharing of assets and the interrelationship created by ADNS there is a greater amount of interdependence than ever before. Although not every casualty can be anticipated there are many likely scenarios for which there should be preplanned responses. Since many situations may not fit the scenarios anticipated there is still no substitute for a detailed understanding of how the systems are configured and what parameters can be manipulated in response to a given set of events.

2. Casualty Conditions

Managing military communications assets is much more complex than their civilian counterparts. Civilian and military both have to deal with natural catastrophes but the military has the added problem of handling the confusion and destruction that can occur in battle. It is critically important for the military to handle problems efficiently and effectively because delay or mistakes could result in loss of life.

The types of casualties that will impact the network or AS are in general the loss of network components. In a stable, shore based network, in any situation other than a major natural disaster, the most likely casualties are failures of individual routers or parts of a link (i.e., a broken cable between two routers). These are also likely occurrences for ADNS nodes as well. However due to the physical grouping of many parts of the

network on the ships in the AS there is the potential for battle damage to remove large sections of a network at one time. As a result the casualty planning process for sea-based ADNS node must necessarily be more detailed than for their shore based counterparts. The types of scenarios that should be anticipated are, in order of increasing complexity: loss of a single radio on one ship, loss of an entire backbone network and loss of an entire ship. Any other scenarios will likely be various combinations of these three.

a. Loss of a Single Radio

This problem includes not just radio failures but a failure of any component in the path to the radio. This includes any of the components from the CAP to the Radio system (see the ADNS block diagram in Appendix A). The magnitude of this problem is inversely proportional to the size of the affected ship (i.e., small ship - bigger problem, big ship – smaller problem). On a smaller ship if this is the only ADNS circuit then the ship has lost connectivity with the backbone.

On a ship that is active on at least two circuits there is the potential for that ship to have established a virtual link to give other platforms access to the backbone. If this is the case then, although the affected ship may still be in communications through its remaining circuits the platforms for which its was relaying information may be cut off. This scenario illustrates the necessity for each platform to understand its role in the AS as a whole.

b. Loss of a Backbone Network

There are two types of links in current ADNS configurations: broadcast and point-to-point. For a point-to-point links, such as Challenge Athena, this failure could occur through a failure at the termination point, the NCTAMS. Since each ship

communicates on these circuits through the NCTAMS then a shore side failure can prevent any ships in the AS from using that network. The loss of a point-to-point link is not as big a problem as the loss of a broadcast network for two reasons. First, from the individual ship's perspective, since these point to point circuits are in general currently available only on larger ships, the impact of a single loss does not have a great impact due to their more extensive radio suites. In addition since the point-to-point links terminate at the NCTAMS which is also likely to be participating on several different subnets the only real loss is network redundancy. Second, from the AS perspective, because it is a point to point link the impact will not be felt as widely through the AS since fewer ships will be using those circuits as compared to some of the broadcast circuits.

For broadcast links the loss of an entire backbone has more serious implications. In satellite based broadcast links the single point of failure is the satellite itself. The loss of a broadcast link could leave large gaps in the connectivity of the AS. Once again for ships operating on a single link this failure can result in lost connectivity to the backbone. For the AS as a whole it can also affect routing protocol overhead. If the lost subnet was a high capacity one the overhead imposed on the remaining lower capacity backbone may be crippling. For more on overhead see the discussion in Appendix B.

Although it is unlikely if the lost broadcast or point-to-point circuit was the sole source of connectivity with the ABR then connectivity with locations outside the AS will have been lost.

c. Loss of a Ship

Besides the catastrophic loss of a ship due to crippling battle damage or sinking a ship may be lost to the AS due to the failure of any of the components in the system that are in the common transmission path. The router and the CRIU are both single point failure items whose loss can remove a ship from ADNS communications. The magnitude of this problem is directly proportional to the size of the ship. Larger ships, such as aircraft carriers, cruisers and command ships, have more extensive communications suites and can be expected to be participating in ADNS on many networks. Their loss will have a much greater impact on the AS than the loss of a smaller ship.

In either case the actions taken in this situation should be focused on determining the overall health of the network. Specific questions to be asked include:

- Is the backbone contiguous? If not, which ships are no longer connected to the network. What subnets are lost if any? Can virtual links be established to restore connectivity to some platforms?
- Was the OTCC on the lost ship? If so who is the backup?
- Can the AS continue to function in this reduced state? Is the overhead on the remaining subnets too high to support passing traffic. If not what actions are required to restore the AS?

3. Specific Considerations for Mission Changes or Casualty Conditions

a. Designate Critical Applications

In heightened DEFCON or EMCON conditions there is a need for traffic control. The transition from peacetime to wartime or hostilities brings with it a shift in

priorities. The mission has changed and the new mission brings with it a need for not only more communications but more rapid communications. Consequently you need to free bandwidth to support the increased level of communications. Stop or limit routine non-mission critical traffic to make way for mission critical information. To do this rapidly and smoothly there must be a plan to restrict access to the network. Access can be restricted either by application or by host or both.

Regardless of the implementation specifics the policy must be promulgated in advance by the OTCC to ensure uniform compliance. ADNS has created a unique situation that can give external communications access to every level in the chain of command. Due to the automatic message handling nature of ADNS there is potential for crewmembers without the appropriate level of situational awareness to send network clogging traffic from one ship that affects the entire AS. This could happen if platforms are permitted to decide independently what applications or hosts to allow access to the network in different situations. Since the process is automatic the policy should do more than inform, it should direct configuration shifts that prevent these unwanted transmissions.

b. OTCC Location

Smaller platforms are, by their lack of redundancy, more vulnerable than larger ships with respect to their connectivity to the backbone. The logical location of the OTCC is on a platform with multiple connections to the AS to avoid isolating him in the event of a casualty. Having the OTCC on a smaller ship could isolate the OTCC more easily because it takes a much lesser magnitude casualty to cause that ship to lose connectivity with the backbone.

c. OSPF Adjustments

(1) Metrics. The most likely reason for changing metrics values is due to a shift in roles among members of the AS. Load sharing among common capacity circuits should be weighed against the need to ensure a minimum performance in support of platforms with a significant mission function. (Casey, July 1997)

(2) Priority. The priority assigned to a host or application will likely need to be adjusted as a result of a mission shift, as opposed to as a result of a casualty. When a new mission begins those applications whose importance to the mission has increased should be assured better level of access than less critical or routine traffic. This is done through a resorting of priorities

The assignment of Priorities and metric values should be optimized to ensure that the right users on the right platforms benefit from the system configuration. For this to be done effectively requires an understanding of both mission requirements and communications capabilities.

(3) Hello Interval. Adjusting the Hello interval can have a significant impact on the overhead imposed on the system by OSPF (see Appendix B.). If, due to mission shifts or casualties, the capacity of least capable platform has been reduced significantly the loading caused by overhead slows the passing of traffic then reducing the Hello Interval should be considered. The negative impact of this is a reduced response time by the network to changes in topology. (Casey, July 1997) But if the alternative is no, or unacceptably slow, communications then it becomes a necessary recovery step. Should the situation improve then restoring or at least reducing the interval in the direction of its original value should be taken as soon as possible.

VI. CONCLUSIONS AND RECOMMENDATIONS

Because the routine operation of ADNS requires little operator intervention the tendency might be to take a more hands off approach to communications management when in fact the opposite is true. ADNS creates a more reliable, efficient and robust communications environment by creating a mobile, radio-WAN interconnecting the Navy's operating forces. To take full advantage of these enhancements requires a network oriented approach to mission planning and execution. With ADNS a network-centric perspective is required of every ship participating in the AS. Each platform is in some way a part of the backbone. Failure of a ship to understand its responsibility with respect to the network as a whole could be disastrous for the mission.

Both (Casey, July 1997) and (Casey and Stell, June 1997) discuss the need for higher level doctrine addressing how best to employ ADNS. This thesis can be used as a starting point for developing a tactical communications management doctrine that can be used by both tactical and communications planners alike when preparing for operations using ADNS. This thesis is also written to provide the at-sea communications managers with information that can be used as a pre-mission tool for developing response plans for various operational conditions. These goals are achieved by:

- Consolidating the information necessary for a management level understanding of the operation of ADNS.
- Highlighting the conceptual difference in our methods of communication as a result of implementing ADNS.

- Providing a consolidated summary of the key elements to be considered when conducting mission planning.
- Providing the ideas to be used in "what if" scenarios by those responsible for managing ADNS systems.

A. AREAS FOR FUTURE RESEARCH

OSPF can support the assignment of up to four metric values. Because no applications currently exist that use more than one, ADNS does not exploit this capability (Casey, July 1997). The ability to assign additional metrics could be used to provide a finer level of control, increasing network efficiency.

The possibility of discarding Hello packets during periods of congestion (Hello Packet Spoofing) is discussed in (Casey, July 1997). The need to reduce OSPF overhead to alleviate congestion warrants additional research into both Hello packet spoofing and Hello Interval adjustment. Operational data on Hello Interval adjustment could be used by operators when deciding what values to use when making Hello Interval adjustments.

ADNS is capable of multicast transmissions via MOSPF. When the same data is passed to multiple platforms multicasting can improve efficiency and reduce the overall amount of traffic on the network when compared to a unicast transmission of the same information. The obstacle blocking the widespread availability of multicast capability in commercial products is conquering the transport protocol problem of providing reliable delivery. Because of the bandwidth limitations inherent in radio systems the use of multicasting in ADNS should be maximized once the reliable multicast problem is overcome.

APPENDIX A. ADNS FUNDAMENTALS

A. INTRODUCTION

1. What is ADNS?

The Navy's Automated Digital Network System (ADNS) provides a means for ship's to centralize and automate the operation of multiple independent radio communications systems into an efficient communications network. ADNS provides connectivity for transmitting bits (which may represent voice, video or data) creating a seamless ship to ship and ship to shore communications network. By managing all of the radio assets within one system, ADNS creates a reliable multiple path communications network. This network is essentially a radio-based Wide Area Network (Radio-WAN) (See Figure A.1). What constitutes the internals of the Radio-WAN are those radio systems configured to support ADNS.

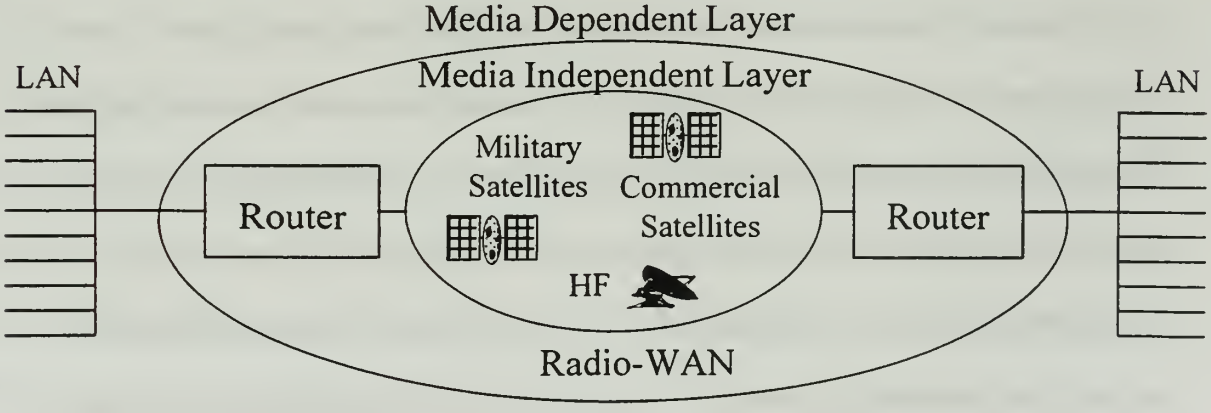

Figure A.1 ADNS as a radio-WAN

Although currently a Navy specific installation, ADNS is like any other LAN/WAN Internet connection utilizing commercial products. Applications need only adhere to the established Internet protocols that ADNS has adopted. This allows a sense of transparency of applications to ADNS. It is also an open-ended system that allows for future expansion. ADNS allows a plug and play like addition of radio links in a process completely transparent to the user.

2. What is ADNS good for?

A group of platforms, linked by ADNS, create a radio-based packet-switched WAN. By using existing Internet technology and open standards users of ADNS have seamless transparent access to the Internet. Using a load balancing concept ADNS spreads traffic equally across the appropriate radio links such that the available capacity is the sum of all the links. ADNS does not provide additional bandwidth instead it multiplexes the bandwidth that is already available.

There has recently been an insatiable demand for Internet access in areas never previously deemed necessary and although Internet technologies are relatively new, limitations are being experienced on traditional wire/fiber transmission paths. The primary purpose of wireless data transfer is for communications with mobile platforms. This capability already exists in various forms. However, ADNS provides a robust means of choosing the most efficient set of paths to transfer data in a way that is transparent to the user. It allows existing stovepipe systems to be integrated into one common data transmission network. When linked with a fixed shore site, to provide wire/fiber connectivity, this network becomes, in essence, a mobile extension of the Internet.

a. Mobile Platforms

Although ADNS was specifically designed for the Navy, it's commercial potential is great. The easiest technology transfer can be applied to maritime platforms. Commercial and research ships have similar needs as the Navy for transferring data to and from shore sites. Imagery (such as weather) transfer, e-mail, Internet access, and file transfer capability are becoming essential tools necessary to accomplish everyday tasks. Commercial aircraft crew and passengers can also benefit from these same capabilities. Cellular phones in automobiles are commonplace. Some cars already receive one way satellite position information using the Global Positioning System (GPS). Currently there is even auto industry research into providing cars with Internet access. The field of mobile communications has become increasingly complex and will continue to grow. However, what should be avoided is a spaghetti-like architecture of different transmission paths linked to different applications.

The traditional way to adopt new data transfer technologies is to implement a stovepipe system with its own dedicated transmission path. Mobile platforms, especially large ones like ships, typically have more than one transmission path for data transfer. However, if data is to be transferred, a dedicated radio link has to be assigned to a specific application. An application can not share different links or be distributed. The same is true for aircraft. Although more limited in space, aircraft too have different radio links which transmit data in a stovepipe fashion. The requirement for data transfer capability in autos is a relatively new concept. However, the reality of cellular phones and GPS combined with the possibility of Internet access already points to multiple transmission paths. Wireless communications do not have to be limited to just

mobile platforms. It can also be a viable alternative to traditional shore links especially if they are saturated or can not be established, for example, in remote areas where the infrastructure just doesn't exist.

b. Alternative to Wire/Fiber Transmission

Traditional shore transmission paths have been saturated with the increasing number of Internet users. Although much research has been done in alternative technologies to alleviate this congestion, such as installing optical fiber, these solutions often require investing in a whole new and different infrastructure. However, ADNS does not provide the same high capacity data transfer capability as shore backbones but instead allows an alternative to traditional mediums for transmitting data without worrying about infrastructure changes. Wireless data transfer could also be an attractive short term solution for areas where the infrastructure doesn't exist or is temporary such as in remote regions. A parallel to this can be seen in many lesser-developed countries where cellular telephones have proliferated because of inadequate landline telephone networks.

3. What Does ADNS Do?

A mobile platform can be thought of as a roaming Local Area Network (LAN). What existed onboard U.S. Navy ships prior to ADNS was a potpourri of different LANs and radio systems. If data was to be transferred to and from a ship, a different radio system was used for each application. ADNS allows platforms with more than one transmission path to integrate these different systems via one black box (ADNS), which then distributes data throughout the different paths in the most efficient manner. This method is desirable for several reasons.

a. Load Sharing

If one or more transmission paths fail or are congested, ADNS can redirect data flow to open channels, which leads to an increased quality of service (QoS). ADNS can distribute data flow much more efficiently than the present stovepipe system. For example, a video teleconference (VTC) often inundates bandwidth, leaving other applications looking for an open transmission path. Other applications such as e-mail can be redirected to less congested channels instead of being stacked in a queue, waiting for transmission.

b. Cost Effective Bandwidth

ADNS can direct data from different applications through desired transmission paths. This can be done to preferentially use the most cost-effective means for data transfer.

c. Leverages the Existing Internet

Another big appeal for ADNS, and one of the main reasons why the Navy has developed it, is that ADNS ties together the existing stovepipe communications architecture. There is no need to create a brand new infrastructure. Existing organizational LANs can be connected to ADNS and have access to the full range of communications assets available to that unit.

d. Flexibility

The use of open protocols and Commercial off the Shelf (COTS) hardware creates a very flexible system. Modifications or additions to the shipboard LAN have no effect on ADNS. By using IP routers as the interface between ADNS and the shipboard

LAN modifications on one side of the router are transparent to the other. Adding a new radio system is not much more complicated than adding a new circuit card.

4. How does ADNS work?

The easiest way to visualize how the system works is through an example. Figure A.2 is a high level block diagram of ADNS and general description of operation.

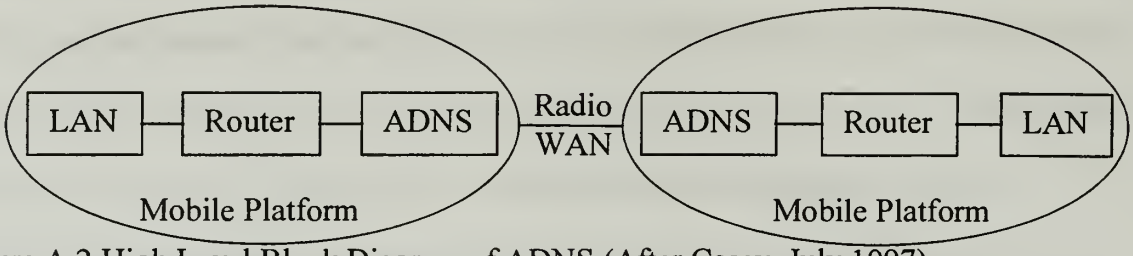

Figure A.2 High Level Block Diagram of ADNS (After Casey, July 1997)

Suppose that a user on a ship at sea wished to transfer a file to another user on a different ship. Let us also assume that both users' computers are connected to their respective shipboard LANs. When the originating user is ready to send the message he simply clicks on the appropriate button to send the message on its way via the ship's LAN.

The size of most data files will necessitate their being broken into multiple IP datagrams. The router, processing each datagram independently, uses the Open Shortest Path First (OSPF) protocol to determine the best path(s) to reach the destination. If there are multiple equal cost paths the router will balance the load amongst them. Similar to a packet switched network a single message may be routed via multiple paths. The router then forwards the datagrams to ADNS. (Casey, July 1997)

ADNS prioritizes, queues and transmits the datagrams on the selected radio system. The transmitted datagrams transit the Radio-WAN much the same way as in a

packet switched network. At the destination there is a mirror image of the transmitting site. Arriving IP datagrams pass back through ADNS to the router and onto the LAN where they are received and reassembled at the destination host. (Casey, July 1997)

This example described the transmission of one message to one destination via a single RF path. To understand the system's true potential, envision multiple ADNS capable platforms communicating simultaneously from multiple applications via multiple RF paths.

5. ADNS Advantages

a. Removing humans from the loop

In current naval communication systems, messages are generated on personal computers or workstations. These messages are transmitted via LAN, (or by use of magnetic media such as floppy disks where no LAN exists), to the communication center. The messages are then processed by technicians and transmitted. This process introduces time delays ranging from minutes to hours. ADNS eliminates the need for human processing of messages by establishing a direct connection from any node on the LAN, through the transmitter, to the receiver at the intended destination. The result is complete automation of the transmission process, with total elimination of any handling delays caused by human interaction.

b. Load Sharing

Most naval vessels maintain at least two operational communication channels at all times. The reason for multiple channels is a legacy one - systems were developed such that only certain types of information could be transmitted and received over each channel. This frequently results in one or more channels being completely

silent, while another is backlogged with traffic. The Load Sharing Feature of ADNS was specifically designed to alleviate these backlogs by making more efficient use of all operational communication channels. This is accomplished assigning a "cost" value to each network. Message queues in each CAP are monitored and messages are routed evenly across equal cost circuits.

c. Optimal use of bandwidth

Network costs are assigned such that higher capacity circuits are assigned lower cost values. ADNS maximizes throughput by finding the lowest cost path for a message to reach its destination. The combination of removing humans from the loop, load sharing and using the lowest cost paths discussed above results in a four-fold increase in throughput during peak traffic times. This is a direct increase in the bottom line throughput of the communications system without purchasing additional transmitters.

d. Communications Agility

ADNS provides the capability for two units that do not share a common communication channel to maintain communications. As long as each unit is operating at least one communication channel and at least one node on the network is operating both channels simultaneously, communications can occur. This process is completely transparent to the users, and occurs with no human intervention. This is analogous to Internet packet delivery. Few end systems share a common communications channel (that is, they are on the same network segment).

e. Transparency of installation and use

The installation of ADNS is totally transparent to the end users. It merely appears that a new router has been added to the LAN with links to many other LANs.

There is no major LAN or transmitter reconfiguration that is required. Additionally, there are no major infrastructure modifications (cooling, ventilation, etc.) required and power requirements are modest.

f. Logistics

The entire installation is small and lightweight, allowing it to be installed in any unused space without impacting shipboard weight and balance.

g. Ease of upgrade

Following initial installation, upgrading of ADNS is quite simple. Addition of new communication channels can be accomplished through the installation of the appropriate CAP cards. Adding capabilities to ADNS itself, such as installing successive builds as they become available, is as simple as downloading the new software. Router reconfiguration is a relatively simple matter as well.

h. Single point for Communications Management

ADNS provides a single point for monitoring all communications, both incoming and outgoing. Prior to ADNS, monitoring all communications was much more difficult due to the lack of interconnection between stovepipe systems. Each of these systems had to be monitored separately. This monitoring capability is available locally via the local net manager's workstation, or remotely from the Network Operations Center.

i. Ability to transmit all types of data

Essentially, ADNS transmits Internet Protocol (IP) datagrams from one router to another. It is the applications on these LANs that decode the datagrams and put the information contained in them to use. Therefore, ADNS can transmit text, graphics,

voice, or video applications over existing channels, without the need for developing expensive new stovepipe systems to support each new application.

6. ADNS Disadvantages

a. Cost of installation

The high initial cost of an ADNS installation is a large obstacle to its widespread use. However, new technology, innovation, and mass production of ADNS should continue to drive costs down. The hardware used in an ADNS installation is COTS equipment but it is very implementation specific. It is unlikely that a unit will already possess equipment that can be modified for ADNS in order to save money on an initial installation. However, future builds of ADNS are planned that will incorporate more readily available hardware. (Casey, July 1997)

B. ADNS OPERATIONAL DESCRIPTION

1. Overview

The behavior of the Radio-WAN created by ADNS is the same as a terrestrial WAN. The router on one platform still "talks" to routers on other platforms, but at a slower rate than if they were connected by wire or fiber. Some of the circuits used in the Navy's ADNS program, such as HF and UHF have transmission rates in the 2.4Kbps range. The insertion of the ADNS hardware and the RF transmission path is simply a conduit for creating a router based network. ADNS deals strictly with IP datagrams. Although some encapsulation occurs as a result of the handling process the underlying

packets are not altered and thus the path between destinations is in essence transparent to the routers.

Figure A.3 below shows the relative position of each component in a typical ADNS setup. The minimum component mix needed for a complete ADNS installation consists of: LAN-Router-CRIU-CAP-Cryptographic Device-Modem-RF System. From the Channel Access Protocol (CAP) to Router Interface Unit (CRIU) back (to the left) there will be only one of each for a given installation. From the CAP forward (to the right) there will be one chain for each radio system that is part of the system (i.e. there may be a UHF SATCOM chain, a UHF LOS chain, an SHF chain, an HF chain, etc.). In this particular configuration there are three RF paths connected to ADNS.

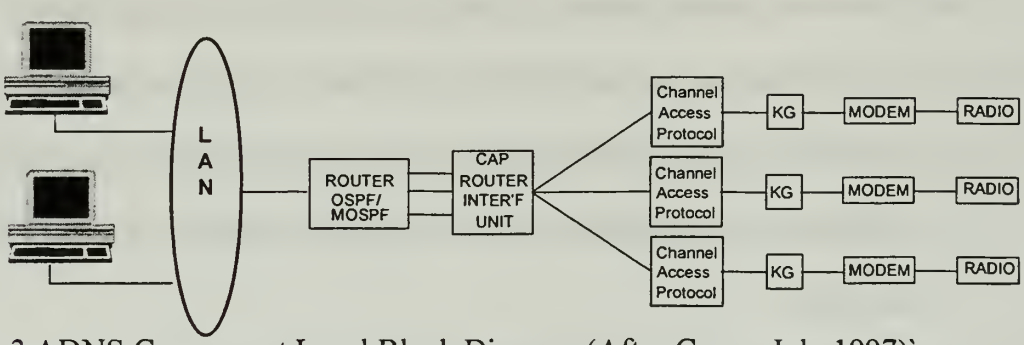

Figure A.3 ADNS Component Level Block Diagram (After Casey, July 1997)`

As discussed earlier, the router accepts outbound datagrams from the LAN and selects the best path for reaching the destination. The CRIU, which interfaces between the router and CAP, assigns a priority to outbound IP datagrams. Priority is inferred based on both the source application (logical port number) and the host (IP address) from which the message originated. At the CAP the message is placed in a queue to await transmission. Messages in the CAP queue are sorted by the priority assigned by the CRIU. (Casey, July 1997)

When the message leaves the CAP it passes through a cryptographic device. The standard Navy ADNS configuration operates at the secret high level of classification, thus all information entering the RF network is link encrypted. This:

- Conforms to existing practice.
- Provides resistance to AS spoofing.
- Provides limited content confidentiality/authenticity protection (because this layer of encryption is stripped off at each routing point). Although this provides protection during transmission it does not provide content security once the information passes through the cryptographic device at the receiving end.
- Provides opportunities for secure tunnels such as Unix Secure Shell (SSH) or Network Encryption System (NES), which deal with IP datagram encapsulation (IP datagrams inside other datagrams). These encapsulated IP datagrams are transmitted by ADNS in the same manner as any other IP datagrams.
- Does not affect applications that offer end-to-end security (e.g. secure e-mail). Similar to secure tunnels, end system encrypted datagrams are unaffected by the presence of ADNS in the system.

After leaving the Cryptographic device the datagram passes through a modem and then enters the transmitter. Once it leaves the ship the message begins traveling via the predetermined path to its destination. Upon arrival at its destination the datagram, traveling through a mirror image of the originating system, terminates at the host specified in the IP header.

2. Network Features

a. Routing Protocols

ADNS uses three different routing protocols. The primary reason for using these algorithms was that the specifications for all three are in the public domain. More detail on the specifics of each of the protocols as they relate to ADNS can be found in Appendix B.

(1) Open Shortest Path First (OSPF)/Multicast OSPF (MOSPF). OSPF is used as the Interior Gateway Protocol (IGP) for routing within an AS. The specification for OSPF Version 2 is contained in Request for Comments (RFC) 2178. It is a dynamic protocol in that each router maintains a continuously updated database containing the status of all other routers in the same system. OSPF uses a lowest cost algorithm to determine the best path to send a message to its destination. Costs are determined based on metrics values assigned to the various transmission paths. (Moy, 1997)

Multicast OSPF (MOSPF) is used for multicast within an AS. The specification for MOSPF is contained in RFC 1584. MOSPF uses the same lowest cost concept as OSPF except the lowest cost is determined with respect to the group. (Moy 1994)

(2) Border Gateway Protocol Version 4 (BGP4). BGP4 is used as the Exterior Gateway Protocol (EGP) for routing between ASs. Specifics for BGP4 can be found in RFC 1771. BGP4 is not as dynamic as OSPF and makes its routing decisions based on predetermined routes. In ADNS, BGP4 will typically reside at the

shore station in a system. Since BGP4 requires a more stable environment than OSPF the shore station is the logical choice. (Rekhter, 1995)

b. Logical Organization

The naming and logical grouping of the elements in an ADNS network are based on the concepts established by the routing protocols used by ADNS.

The basic unit of an OSPF network is an area. For ADNS a ship is typically considered an area. Certain shore installations will also be areas since the ships need an interface point with other shore based establishments. A number of ships grouped together using OSPF create an Autonomous System (AS). A typical AS consists of a group of Navy ships with some logical connection, such as a common mission. A Battle Group is a typical AS. The emphasis in AS establishment is on mission and not location. The units do not have to be in the same geographic region to be in the same AS. At least one and possibly two or more shore communications establishments will also be a part of an AS to act as the gateway to other Navy networks such as the SIPRNET (Secret IP Router Network) or the Internet. (Casey and Stell, June 1997)

The combined network of RF systems creates the subnet backbone of the AS. Each subnet is a different RF system such as UHF Satcom, SHF Satcom or INMARSAT B. The router on each ship that interfaces with ADNS is established as an Area Border Router (ABR). Each ABR operates OSPF. Part of the data that is maintained in the OSPF routing tables are metrics for each subnet in the AS. In current ADNS installations, metrics values are assigned based on subnet capacity or bandwidth. Higher capacity subnets are assigned lower metric values. The values chosen for these metrics determine how the system performs load balancing and load sharing, as discussed

below. Obviously since each router must maintain a dynamically updated table of every other router in the AS there is a limit to the number of routers which can be managed effectively. This is what drives the upper limit to the size of an AS. (Casey and Stell, June 1997)

The router that acts as the gateway between an AS and other ASs, WANs, or the Internet uses BGP4. The shore establishment usually performs this function since BGP4 needs a stable environment. The OSPF to BGP4 transition acts to hide the internals of the AS from the outside. Routers outside the AS don't need to know the specifics of all the routers inside the AS. They only need to know where the BGP4 gateway into the AS is. Changing missions will prompt changes to an AS. Ships may need to transfer from one AS to another to support operational or training objectives. This dynamic reorganization requirement reinforces the need to shield the internal routing issues of each AS from the outside. Figure A.4 shows the relationship between routers within a simple Autonomous System. (Casey and Stell, June 1997)

The third party routing feature of this type of network is illustrated in Figure A.5 below. If the originating and destination ships are not operating a common circuit ADNS will route traffic through a third platform which has connectivity on both source and destination circuits. By maintaining the status of other ships in the AS, ADNS can determine the best path to ensure delivery of each message. The diagram shows how the sending ship's router (R1) will send via either EHF or UHF (or both, depending on the metric values assigned to each RF path) to R3. R3 will then forward via HF to the destination ship, R2.

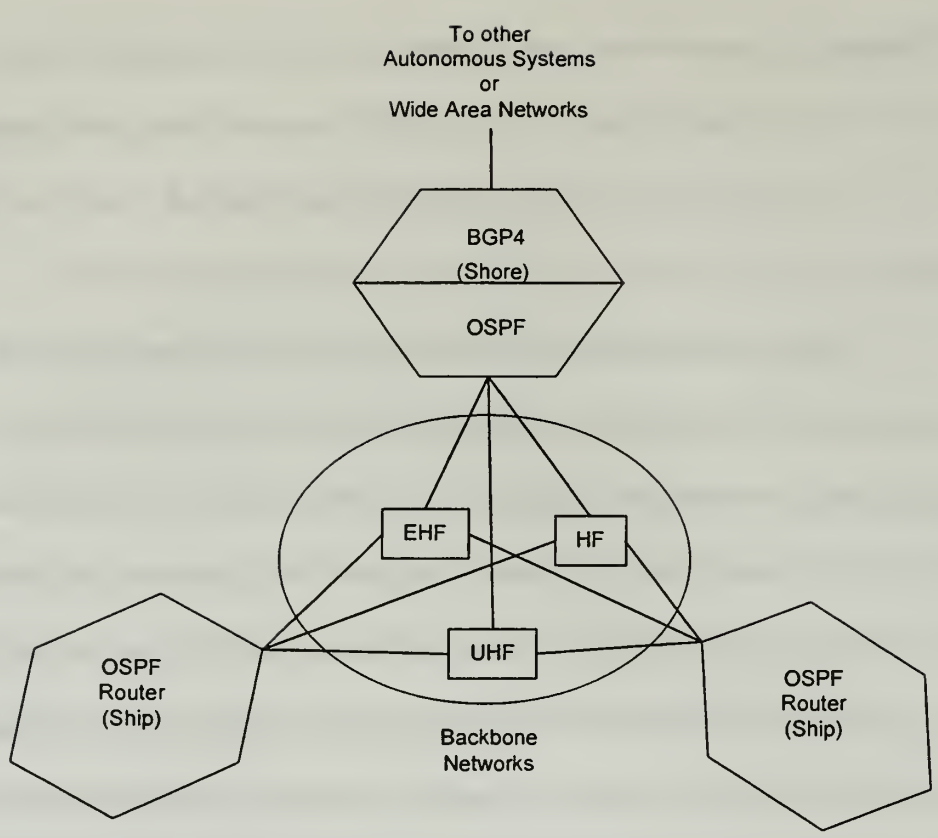

Figure A.4 Relationship Between Routers in ADNS (After Casey, July 1997)

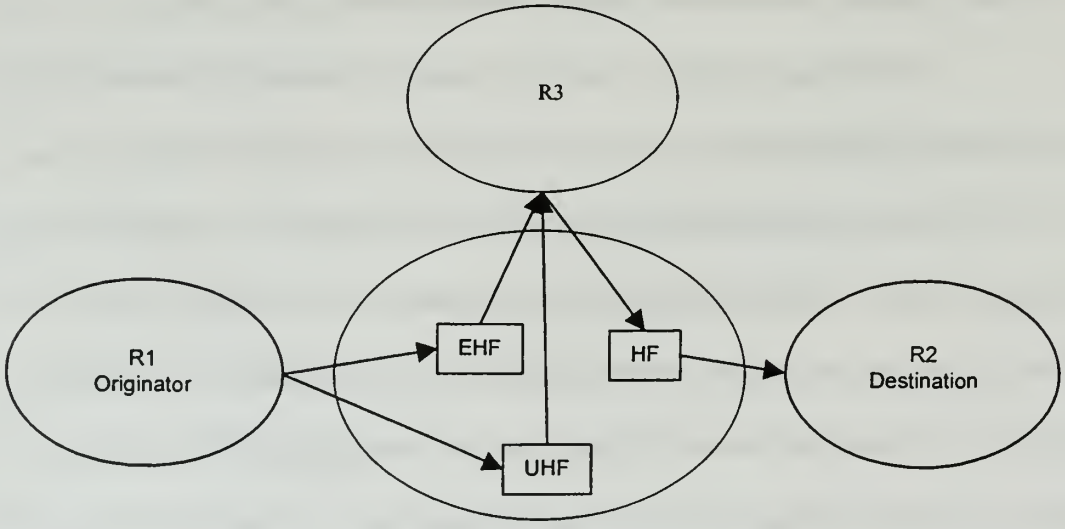

Figure A.5 Third Party Relay Function (After Casey, July 1997)

3. Key Features/Functions of ADNS

The concepts presented in this section have been condensed from (Casey and Stell, June 1997).

a. Priority

Several different methods for assigning priority to outgoing messages were evaluated during the ADNS implementation process. One obvious method, using the built-in precedence field in the IP header, was briefly considered. This idea was quickly discarded since no relevant applications currently use this feature of the IP header. Eventually, a priority scheme was implemented which assigned priorities of 0 (lowest) to 15 (highest). The two methods which proved most useful for assigning priority were based on source IP address (Host), or port number (Application).

This approach has the same advantages and drawbacks of a firewall that uses the same data to make its filtering decisions. The advantage is its practicality. The disadvantage is that it's rather crude and, at the moment requires manual configuration of the router's routing table.

(1) Priority Tables. The CRIU maintains two priority tables. The Source IP table contains the IP addresses of hosts on the associated LAN and the priority which they have been assigned. There are no default settings for this table. If a host is to have an associated priority it must be entered into the table. This table is filled in manually by the local ADNS Manager during initial system configuration and can be updated at any time. The Source IP table contains space for up to 40 entries.

The second table maintained by the CRIU is the Port priority table. It contains the dedicated port numbers used by certain applications and the priority that

has been assigned to that particular application. Just as with the Source IP table above, there are no default values, priorities must be entered manually, and it contains space for up to 40 entries.

(2) Determining Message Priority. The CRIU receives datagrams from the router. The CRIU determines the port number and originating IP address for each datagram and assigns priority based on entries in the Source IP and Port priority tables. Here, a conflict may arise. If the Source IP priority table assigns a certain priority to a particular datagram and the Port priority table indicates a different priority for the same datagram, priority assignment will be made on the basis of Source IP address. This allows priority based primarily on host, and secondarily on application should the host have no assigned priority. If neither the host nor the application have an assigned priority, the CRIU assigns a default value of priority 4. Once assigned, the priority is placed in the IP datagram header and the entire IP datagram is passed to the CAP.

(3) Message Transmission. Following assignment of priority, the IP datagram is forwarded to the appropriate CAP, where it is entered into one of 16 queues based on priority. Datagrams are assembled into transmission units, each of which can contain up to 64 IP datagrams. The size of the transmission unit depends on the capacity of the link. Lower capacity links will have to utilize lower transmission unit sizes. The CAP builds a transmission unit by removing datagrams from the queues in order of priority. Datagrams are removed from the highest priority queue first, until it is empty. Datagrams are removed in sequence, continuing down the priority queues until

the transmission unit is complete or all queues are empty. The transmission unit is sent from the CAP to the corresponding RF transmitter and the process is repeated.

b. Load Balancing

Load balancing is the sharing of transmission load equally among different subnets. When the router selects a transmission path it does so based on the metrics assigned to that RF system. OSPF metrics are based on link capacity, with links having similar capacity being assigned identical metric values. If multiple CAPs have the same metrics values then the router will balance the load evenly by alternating between those CAPs. For load balancing to work effectively the sharing must be done among systems of equivalent capacity. Consequently, when assigning metric values to RF systems it is important that only networks of like capacity be assigned the same values. For example, if a ship is operating two active subnets, HF (which operates at about 2.4Kbps) and SHF (which operates at about 64Kbps) assigning the same metric values to each would overload the HF circuit. The router would divide the load equally between the two, not proportionally. During periods of high traffic density the SHF link could handle the load more effectively than the HF link, which would become backlogged with data.

c. Congestion Control

As described above, each CAP maintains separate queues for each priority (0-15). Should one of these queues become full, the CAP does not provide any overflow queue so additional datagrams with this same priority will be dropped. In order to prevent this situation from occurring, the CRIU monitors the CAP queues and either starts load sharing or issues a Source Quench command.

Each queue in a CAP is allocated a certain queue size to store IP datagrams prior to transmission. The CAP manages this queue space. The CRIU sets a queue threshold, slightly smaller than the queue size, to use as a benchmark to determine if congestion of the queue exists. The gap between the queue threshold and the maximum queue size provides a buffer to allow action to be taken before the queue becomes full and datagrams start being discarded. These queue thresholds are pre-determined and entered into the CRIU by the local ADNS Manager. The congestion identification function operates in the following sequence. The CAP generates a queue report, at intervals specified by the queue report threshold. This report captures the actual queue levels and sends them to the CRIU. These levels are compared to the queue threshold for each queue. If any queue level is greater than the queue threshold, then a congestion condition exists in that queue. The macro behavior of this arrangement is very similar to congested routers in a conventional Internet so TCP, including the Karn and Nagel algorithms, will work without change.

(1) Load Sharing. One of the key features of ADNS is its ability to share the traffic load over available subnets. In current Navy circuits a situation frequently occurs in which one communication channel is overloaded while another is completely idle. The load sharing feature of ADNS alleviates this problem by shifting some of the congestion to the idle channel, thereby increasing throughput and shortening communication system delays. This differs from load balancing in that balancing distributes traffic over channels with similar metric values before congestion occurs. Sharing distributes traffic over similar cost channels because a congestion condition exists.

(a) Restrictions. There are two restrictions on the use of load sharing. First, the traffic being shifted to an alternate channel must be unicast traffic only. Multicast applications introduce a level of complexity that causes diminished returns, making it not worth the effort to attempt to load share using multicast applications. Second, load sharing is only feasible between subnets whose bandwidths are in the same range, meaning they share a similar time delay. Thus, possible opportunities for a load sharing situation are between UHF and EHF, or between SHF and Challenge Athena.

(b) Implementation. The load sharing process begins when the CRIU determines that a congestion condition exists on a subnet in one of its associated CAPs. The CRIU then scans all other compatible (those with similar delay times) subnets to determine if a path from origin to destination exists. If another subnet does exist with a path from origin to destination and no congestion condition exists on this subnet, load sharing commences.

(2) Source Quench. When congestion is determined to exist in the CAP queue for priority n, the CRIU issues a Source Quench ICMP command. This command stops the generation of message packets for all applications and hosts with priority n or less. Assuming compliant TCPs this Source Quench command has been pre-set to remain in force for five seconds. At the end of five seconds, transmission from the affected hosts and applications resumes automatically unless or until another Source Quench command is issued. It should be noted that all applications and hosts require some sort of flow control to ensure that during Source Quench conditions, packets are not discarded but rather stored for transmission when the Source Quench has timed-out.

d. Transmission Control Protocol (TCP) Duplicate Packet Transmission Problems

One of the major early setbacks to implementing the ADNS architecture was solving the problem of TCP duplicate transmissions when initially establishing a TCP connection. ADNS causes the LAN gateway router to act as if it is hard-wired to other routers on other LANs. Thus the router expects to encounter minimal delays (less than 0.5 seconds) in receiving acknowledgments to its TCP packets being sent. In reality, these TCP packets are being transmitted over RF links to distant LANs. The minimum acknowledgment time for a 1500 byte packet over a 2400 BPS connection is in the neighborhood of 5 seconds. When TCP hasn't received packet acknowledgment after 0.5 seconds, it re-transmits the packet. If acknowledgement is still not received after an additional 1 second, TCP retransmits the packet again, and again after 2 seconds, 4 seconds, 8 seconds, and so on. Under optimal conditions, a 1500 byte packet will be sent 4 times over a 2400 BPS connection. The end result is the use of 6000 bytes to transmit 1500, an efficiency of 25%.

(1) TCP Duplicate Packet Rejection. A practical solution, and the one implemented in ADNS, is to design the CRIU to discard duplicate TCP packets before they are transmitted over the RF link. This is accomplished by the use of a table for each subnet that contains the TCP sequence number and time-stamp indicating when the packet was received by the CRIU for transmitting. A TCP original packet and each duplicate packet sent will have the same TCP sequence number. When a TCP packet is received by the CRIU for transmitting, its TCP sequence number is examined. If this number already exists in the table, the packet is rejected. If this number does not exist in the table, it is added to the table along with its time-stamp, and the packet is passed along

for transmission. Each subnet is assigned a TCP duplicate rejection time. If a TCP sequence number has been in the table for longer than the TCP duplicate rejection time, it is deleted from the table. The TCP duplicate rejection time has a default value of 10 seconds. This provides for transmission of the original TCP packet followed by a 10 second delay for acknowledgment. If none is received, the packet is allowed to be retransmitted followed by another 10 second delay. This time delay can be modified by the Local ADNS Manager, based on the latency of the link, for optimum performance.

4. ADNS Integrated Network Management

a. Overview

Network management of ADNS is based on SNMPv1 standards. There are no proprietary Navy protocols to confront, thus allowing the use of standard network management tools and practices. Most of the objects to be managed (hosts, routers, etc) will have agents attached and MIBs will be written for any unique objects. The Navy will adopt a standard, commercial Network Management System (NMS) to provide the foundation for network management. However, there are Navy-specific concerns, such as command and control relationships, which impact network management. For these special requirements, the Navy will create special applications and concepts to the NMS. This section gives a broad description of how the Navy intends to manage ADNS.

Network management of naval nodes is similar to managing shore-based nodes. The fundamental concepts are the same. However, the mobile nature of the nodes makes managing shipboard nodes more difficult. The fact that they are warships makes management more important. Just as there is a military hierarchy there is one for network management in ADNS, where each level has different responsibilities. Network

management is a vital portion of ADNS because the consequences of system errors or failures can directly affect combat effectiveness.

Integrated network management describes how the Navy will manage networks on a distributed basis all the way down to individual objects. They include, but are not limited to: general monitoring, statistic collection, status monitoring, traffic monitoring, trend analysis, network loading, network optimization, configuration control, system configuration, maintenance, problem identification, problem reporting, trouble documentation, system administration, and emissions control [INM Technical Approach].

Network management of ADNS contains three different levels: the Local Control Center (LCC), Autonomous System Control Center (ASCC), and the Navy Operations Center (NOC). The LCC will be responsible for networks at the local level, e.g. within an area (usually a ship). The ASCC will be in charge of networks on a regional level, having several subordinate Autonomous Systems. The NOC will be responsible for all ASCCs in a certain geographic area. This arrangement is consistent with the Navy's organization and its doctrine regarding distribution of authority.

(1) Local Control Center (LCC). The LCC is the network management center at every unit level. There is a local responsibility to monitor and maintain the status of all subnets at that unit. There are three components of an LCC: a Network Manager, Distributed Manager and a Communication Automation Manager.

(a) Network Manager. The Network Manager is network management system software that is obtained commercially. The purpose of the network manager is basically to give the status of the network and individual objects. An example is the popular HP Open View Network Node Manager product (OV-NNM) which has

been in the Navy Tactical Advanced Computer (TAC) contracts since 1991. It provides a topological map representation of a unit's network and shows the status of each object with the use of colors and shapes. However, human interaction is required to interface with the ASCC and the NOC for troubleshooting or maintenance. The specific functions of a Network Manager will be:

- Human machine interface
- Performance management
- Fault management
- Accounting management
- Security management
- Configuration management

The Network Manager will be used as the foundation for the Navy's Integrated Network Management System, where specific applications can then be added on to provide other management functions.

(b) Distributed Manager. Distributed Management is an application that determines what is to be reported locally and what is to be reported to the ASCC and NOC. The Distributed Manager has two mechanisms for discovering if any conditions exist that meet the criteria of its policy rules:

- Notification from the Network manager
- Query from Distributed Manager to Network Manager

The specific functions of the Distributed Manager will be:

- Interpretation and implementation of policy
- Filtering of management information

Although commercial products can provide these functions, the distributed manager in the Navy context specifically describes the policy rules for the communication relationships between the LCC and ASCC.

(c) Communication Automation Manager (CAM). The Communication Automation Manager is in charge of the physical communication hardware and their related requirements. On a ship, they are functions typically related to the radio room. Duties include a communication plan implementation, circuit building, and circuit management. Three areas make up the Communication Automation Manager: the Communication Manager, Site Manager, and Equipment Manager. The specific functions of the Communication Automation Manager will be:

- Security management
- Log Control
- Alarm reporting
- Summarization
- Attributes for representing relationships
- Objects and attributes for access control
- Usage Metering
- Test Management
- Event Report Management
- State Management
- Security alarm reporting
- Object management
- Bandwidth management

- Communication plan management
- Equipment control
- Site configuration management

The Navy specific application for these functions is the use of a remote management tool called the Communications Plan (COMMPLAN). The COMMPLAN will used to direct certain network management functions as described above. This is still mainly accomplished manually by a technician after receiving the COMMPLAN via hardcopy message. However ADNS will allow many of these requirements to be accomplished remotely and automatically via the COMMPLAN transmitted to the Communication Automation Manager. This concept can be applied to commercial industries where it is not cost effective to have the necessary network management expertise at every local site but can instead be centralized at one remote center.

(2) Autonomous System Control Center (ASCC). An ASCC monitors the operation of several LCCs. The Navy's configuration will use its regional shore communications stations, Naval Computer and Telecommunications Area Master Stations (NCTAMS) as ASCCs. The ASCC will receive summary reports from subordinate LCCs. The exact nature of reporting from an LCC to an ASCC is still to be determined but will contain mission relevant information. Such reporting requirements can include:

- Readiness of communication to support the mission.
- Status of communication services.
- Status of hardware and software.

- Information about usage and reliability.

ASCCs can also give direction to LCCs regarding communications posture. This could include such items as prioritization of resources or equipment configuration changes.

(3) Network Operations Center (NOC). The NOC is the next level above an ASCC for reporting network management information. The NOC would basically monitor all nodes in a certain geographic location. For example the Navy has established a NOC in the Pacific and Atlantic regions. Although capable of monitoring detailed network management information, a NOC would be more interested on the overall status of ASCCs and LCCs.

b. Network Management Tools

To achieve the above network management requirements, a vast array of tools are available to all levels of management and maintenance personnel. However, each tool comes with their own training requirement. Therefore the total cost of ownership must be taken into consideration against their utility. The basic tool for monitoring the network is commercially available Network Management System software. Another tool available for the goal of transparent and affordable network management is software that is capable of remote monitoring and maintenance. These can also be available commercially or can be developed to be mission specific. There are always emerging tools on the horizon for new technologies. However, one of the primary reasons why network management techniques lag behind new network technologies is that time is needed to see which technologies will become established as industry standards. ADNS will manage objects primarily through SNMPv1 standards.

That is not to say that ADNS can not adapt any emerging technologies that become industry standards, such as SNMPv2. However, SNMP has proved that it will be around for a long time.

(1) Network Management System Software (NMS). A commercial Network Management System software has been adopted for the foundation of the INM. Network Management System software allows for the basic functions of monitoring nodes and network status. As described earlier, many different types of enterprise management software are available commercially, such as the popular HP Open View Network Node Manager (OV-NNM). Although commercial software provides excellent monitoring tools, proprietary software is often required to achieve other network management requirements. Commercial Network Management System software offers a fairly inexpensive solution that provides a solid foundation of network management tools. Additionally, to provide the flexibility desired throughout ADNS a COTS product is appropriate.

(2) Third Party Applications. An attractive feature of a Network Management System such as OV-NNM is that third party applications can be integrated into it. Especially for organizations like the Navy, solutions to mission specific requirements can not be obtained off the shelf. These mission specific add-ons must be developed independently and then integrated into the existing NMS. Proprietary equipment also requires some kind of integration with the NMS. Such things as configuration management software for specific objects must be obtained from the vendor. For example, companies offer software that can be integrated with an NMS to allow managers to remotely configure their hardware. Third party applications offer

remote management capability. This is the whole purpose of enterprise management. It is very cost effective to centrally manage nodes rather than paying for the necessary expertise at every local level. Although there needs to be some human interaction at every level, full management capabilities are not required down to the local level.

ADNS is a good example of the need for remote management. Implementation of remote management over ADNS will allow managers to configure and manage mobile platforms from a central management location. This, in turn, allows the assignment of minimal personnel at the local level, thus saving on personnel costs. With such standards as RMON and SNMPv2, remote managers can access remote networks in a secure manner and troubleshoot or reconfigure the network. For example, if one transmission path fails, a remote manager can gain access to the system via a second transmission path and troubleshoot the system. The use of more than one transmission path allows the ability to continually manage LCCs and even ASCCs remotely through just one open path. Although ADNS has not adopted such standards as RMON or SNMPv2 yet, the technologies currently exist and can be readily integrated into ADNS.

C. HARDWARE

1. LAN

The LAN will typically be the existing shipboard Ethernet or FDDI network. Hosts on the network will run a wide variety of applications.

2. Router

The router is an IP router that acts as a gateway to the ADNS network. The router can be any commercial router capable of running OSPF. Currently the ADNS program uses the CNX 600 Proteon router.

3. CRIU (Channel Access Protocol to Router Interface Unit)

The CRIU is implemented on a single board computer installed in a VME chassis.

4. CAP (Channel Access Protocol)

A CAP is also implemented on a single board computer mounted in the same VME chassis as the CRIU.

5. Cryptographic Device

Navy ADNS installations use the KG-84 for link encryption.

6. Modem

For each CAP there is a corresponding Modem that performs the analog to digital (inbound) or digital to analog (outbound) conversion of data passing through ADNS.

7. Connectivity Media

Each RF system (e.g. UHF Satcom, EHF Satcom or INMARSAT B) constitutes one network when considering all assets in one ADNS Autonomous System.

APPENDIX B. APPLICABLE ROUTING PROTOCOL CONCEPTS

ADNS uses three open standard Internet protocols to accomplish its routing functions: Open Shortest Path First (OSPF), Multicast OSPF (MOSPF) and Border Gateway Protocol Version 4 (BGP4).

A. DEFINITIONS

The following general definitions are applicable to all three protocols.

- Autonomous System (AS) - "A group of routers exchanging routing information via a common routing protocol (Moy, 1997).
- AS Boundary Router (ASBR) - A router which links an AS to other ASs. (Moy, 1997).
- Interior Gateway Protocol (IGP) - "The routing protocol spoken by the routers belonging to an AS" (Moy, 1997). Although different ASs may be using different IGPs, each AS only uses one. OSPF is an IGP. All ADNS ASs use OSPF.
- Area - A group of networks whose topology is hidden from the rest of the AS. "An area is a generalization of an IP subnetted network" (Moy, 1997). In ADNS each ADNS installation (ship or shore site) will usually be considered an area (Johnson, 1997).
- Backbone - The common area through which areas are attached (Johnson, 1997).
- Area Border Router (ABR) - A router attached to more than one area (Moy, 1997). In ADNS installations it is the area router attached to the backbone (Johnson, 1997).

- Exterior Gateway Protocol (EGP) - A routing protocol used to communicate between ASs. BGP4 is an EGP. Routing between ASs in ADNS is done via BGP4.

B. INTERNAL ROUTING

1. OSPF

a. General

OSPF is a dynamic routing protocol used to communicate between routers in an AS. OSPF is connectionless, operating at the network layer of the OSI model. Each IP datagram is independently routed to its destination based on the destination IP address in the packet header. The full specification for OSPF Version 2 can be found in (Moy, 1997). Except where specific reference is made to the ADNS implementation of OSPF, this description is a consolidation of relevant sections of that RFC.

The dynamic feature of OSPF means that each router maintains a frequently updated link-state database containing information about all other routers in the AS. This information is used to create a table of paths to every other router and network in the AS. Each path has an associated cost. The route by which each packet is sent is the lowest cost path chosen by the router. Costs are calculated based on a dimensionless metric value assigned to each path.

OSPF allows for the subdivision of an AS into areas to reduce the communications required to maintain the status of the network. When areas are established the topology within an area is hidden from the rest of the AS and the topology of the rest of the AS is hidden from that area. In an AS that has not been divided into

areas each router has an identical link state database. When areas are used only those routers connected to the same area have identical databases.

It is the job of the ABR to represent the consolidated route structure of the backbone into its area and to provide the rest of the ABRs in the AS with the information necessary to route information into its area. To perform this function the ABR runs a copy of the algorithm for each area to which it is attached.

When areas are used the backbone is also considered an area. It contains all ABRs in the AS. "The backbone must be contiguous. However it need not be physically contiguous; backbone connectivity can be established/maintained through the establishment of virtual links" (Moy, 1997). A virtual link is established by configuring one area to act as a relay for another area. For example, area A is connected to both the backbone and area B. Area B is only connected to area A. Area A can be configured to act as a virtual link to connect B to the backbone. The route to B is advertised through A.

b. The Link State Database and Routing Table

Each router on the network maintains a link state database that includes the cost for each connection in the network. Since the costs associated with a given connection are direction sensitive the database contains both a "to" and "from" entry for each connected network or router. For example, if two routers, R1 and R2, are connected there will be entries for R1 to R2 and R2 to R1. The cost for each may be different, depending on the metric values assigned.

The router calculates a routing table of shortest paths to each destination from the link-state database. This table has three columns: destination, next hop and distance. There is a line item for each network. Although the algorithm calculates the

entire path, only the next router (next hop) is entered in the routing table. Distance is the total cost to the destination network as calculated from a particular router. Since the shortest route to any destination depends on the starting point, the routing table will be different for each router.

c. Link State Advertisements

The Link State Database is built from the information provided in Link State Advertisements (LSAs) received by the router. LSAs describe the current state of the connections within a network as seen by a given router at a specific time. There are five different types of LSAs:

- Type 1: Router-LSA. Describe the links a router has to an attached area. Included in this description is the metric value assigned to each link.
- Type 2: Network-LSA. Sent by the DR on Broadcast and NBMA networks this LSA lists all routers connected to the network.
- Types 3 and 4: Summary-LSA. There are two types of Summary-LSA. Sent by an ABR this LSA describes a route to a destination outside of that area but still inside the AS. One type gives routes to ASBRs. The other type gives routes to networks. Included in the Summary-LSA is the metric value for the entire route to the destination.
- Type 5: AS-external-LSA. Sent by an ASBR this LSA describes a route to a destination in another AS. This LSA also contains a metric value describing the cost of the route.

d. Routing Protocol Types

To establish and maintain the status of the network information in various

forms must be passed among routers in an AS. To accomplish this OPSF uses five different protocol packet types: Hello, Database Description, Link State Request, Link State Update and Link State Ack. With the exception of Hello packets these packets are sent only over adjacencies. Among the information found in each packet is:

- Router ID. Uniquely identifies the originating router.
- Area ID. Identifies the area to which the originating router is connected and which is the source of the packet. Packets are associated with areas vice routers since routers can interface with more than one area but the information in a packet describes relationships with respect to an area.
- Authentication. Each packet is authenticated, thus only trusted routers may participate in a network.

The Hello packet is used to find and maintain neighboring routers. It is also used in the Designated Router (DR) election process. Among the additional information included in a Hello packet is:

- HelloInterval. Interval at which Hello packets will be generated. This value must be the same for every router on the network.
- RouterDeadInterval. Elapsed time from receipt of last Hello packet before a router is declared down. This value must be the same for every router on the network.
- Designated Router. IP address of the DR. If no DR has been elected this field is set to 0.0.0.0.
- Backup Designated Router (BDR). IP address of the BDR.

- Neighbor. Router ID of any router whose Hello packets have been received by the originating router within the last RouterDeadInterval seconds. This field is repeated as necessary, once for each neighbor.

The Database Description packet is used between two routers when adjacency is being established. Information in the packet includes:

- DD Sequence Number. Each packet is sequentially number to ensure continuity between the two routers exchanging data.
- LSA Header. The header information (vice the fully database entry) for each LSA in the database. Due to packet size limitations each packet can only hold a finite number of LSA headers. Consequently to fully describe a database will usually require multiple Database Description packets.

Generated in response to a Database Description packet, the Link State Request packet is used to request missing parts of a link state database. The Link state request identifies the LSA for which an update is needed. Each Link State Request can request multiple LSAs. Similar to a Database Description packet the packet can contain multiple LSA header fields.

Link State Update packets are sent in response to Link State Requests or when the status of a router changes. In addition to the LSA header the packet also contains the full LSA. Each packet can contain multiple LSAs and they can have originated from different routers.

Link State Acknowledgement packets are sent to acknowledge receipt of Link State Updates. The body of the Link State Acknowledgement packet lists the LSA headers for which receipt is being acknowledged.

e. Establishing a Connection

To support the dynamic nature of the protocol OSPF routers must communicate often to pass information regarding the status of the network. The functions performed by a router when it is first brought into the network can be divided into a sequence of four steps; discovering neighbors, verifying two-way communications, electing a designated router (for broadcast and non-broadcast multi-access (NBMA) networks) and, if appropriate, establishing adjacency.

(1) Discovering Neighbors and Verifying Two-way Communications. To ensure delivery of data each router in the AS must have an accurate picture of the current state of the network. The first step in forming this picture is to determine what other routers are available. This process of neighbor discovery is accomplished using the Hello Protocol. Each router will upon startup and periodically thereafter send Hello packets to other routers in the AS. The Hello packet allows each router to advertise its status to other routers.

The hello packet sent by a given router contains an entry for every other router for which it has received a current hello packet. As the newly started router receives Hello packets from other routers it updates its own Hello packets. At the same time other routers in the network are updating their packets by adding the new router. Two-way communications are verified when a router see itself listed in the Hello packet of another router.

(2) Electing the Designated Router (DR) and Establishing Adjacency. On networks with multiple routers (broadcast and NBMA networks) maintaining an updated network status on all participating routers can contribute a

significant amount to the traffic on the network. To help control the amount of traffic on these types of networks the OSPF protocol provides for the electing of a designated router and the establishing of adjacencies. To minimize traffic only adjacent routers exchange routing information updates.

Each router is assigned a router priority. That priority is included as a data field in the Hello packet. The designated router is usually the router with the highest router priority. When the new router enters the network it looks for a DR. This discovery process is done by the examination of incoming Hello packets. The hello packet generated by each router indicates which routers it thinks are the DR and Backup DR (BDR). If a DR has not been elected and the new router has the highest priority in the network then it will become the designated router. If there is already a DR then the new router will accept the existing DR, even if the new router has a higher priority. Although it makes it harder to identify which is the DR, this method creates less disruption for the network since shifting of DRs requires updating the databases on all routers in the network. This disruption could cause delays in routing of data on the network while the router databases are being updated.

In addition to the DR there may also be a BDR. This is to avoid network disruption when the DR fails. Since each router already knows the identity of the DR and BDR the shift to the BDR on a loss of the DR will not require excessive network communications to reestablish the state of the network. To minimize the number of shifts the most dependable router in the network should have the highest priority so that it will eventually become the DR.

Once the DR and BDR have been elected the process of forming

adjacencies begins. Not all routers become adjacent. Routers only become adjacent to the DR and BDR. To become adjacent means that the link state databases of the two routers are synchronized. To synchronize databases the routers must exchange database status information. This is done via Database Description packets. The two routers establish a master-slave relationship for this Database Exchange Process. The master sends the status of its database via Database Description packets. The slave receives these packets and acknowledges receipt by sending a Database Description packet with the same DD sequence number and its version of the LSA header information back to the master. Each router then compares the LSA information to its own database. If either router has data that is older than the other router's it requests an update via a Link State Request. When the Database Exchange Process is complete both databases are identical and are considered synchronized and the routers are considered to be adjacent.

f. Network Maintenance

To ease the communication overhead associated with maintaining the network several of the OSPF protocol packet types can be sent via IP multicast. There are two IP multicast addresses used in OSPF, AllSPFRouters and AllDRouters. All routers running OSPF should be configured to receive packets addressed to AllSPFRouters. Each router sends Hello packets using AllSPFRouters. The DR will also use AllSPFRouters when sending Link State Update messages to all adjacent routers. Adjacent routers use AllDRouters to send Link State Updates to the DR and BDR.

It is important to note that since it is only one hop from the DR or BDR to any adjacent router then all of the packets that travel only over adjacencies travel only one hop. Since Hello packets are sent to immediate neighbors this means that no OSPF

packet is required to travel farther than one hop from its source. The only exception is for virtual links that may need to forward packets to their ultimate destination.

Maintaining the status of the network current requires the periodic passing of all of the different types of messages at varying intervals. Hello packets are sent at an operator selectable interval set by the HelloInterval setting in the Hello packet. The value chosen should be significantly less than the RouterDeadInterval to avoid unnecessarily terminating connections. Database Description packets are retransmitted by the DR at fixed 30 minute intervals. Link State Requests, Updates and Acknowledgements are sent as needed in response to changes in the network topology.

g. Packet Routing

Routing of packets is done in three steps. Intra-area routing through the area of the originating network, inter-area routing across the backbone area and intra-area routing through the area containing the destination network. The algorithm finds the combined set of paths with the smallest cost. The router consults the routing table for the destination address of each packet and forwards it to the Next Hop router listed in the table. The process is repeated at each router until the destination is reached.

2. MOSPF

a. General

Multicast OSPF is an enhancement to the OSPF routing protocol that allows for the multicasting of IP datagrams (Moy, 1994). The full specification for MOSPF can be found in (Moy, 1994). Because it relies heavily on the existing OSPF structure this discussion of MOSPF serves to highlight the important differences between the two protocols. This description is a consolidation of relevant sections of the MOSPF

RFC.

b. Characteristics of MOSPF

MOSPF adds one additional LSA to those already used by OSPF. The group-membership-LSA serves to identify multicast group members in the existing OSPF database. Much like OSPF the multicast extension calculates a shortest path tree for transmitting datagrams, using the same metric values as OSPF. However, unlike OSPF, this tree is calculated on demand, when the first datagram in the transmission is received.

MOSPF also differs from OSPF in that in OSPF IP datagrams are routed based on destination IP address only, in MOSPF datagrams are routed based on both source and destination addresses. When routing datagrams MOSPF will take advantage of any common paths among the destination addressees. The datagram will not be replicated until the paths diverge.

MOSPF does not allow for equal cost multi-path routing. Only one path will be selected for each destination IP address. Due to the division of an AS into areas each router does not have a complete picture of the AS since only summary information is advertised across area boundaries. As a result, the routing of datagrams may be less efficient due to the hiding of paths performed by the ABRs.

C. EXTERNAL ROUTING

1. BGP4

a. General

BGP4 is a routing protocol for use between autonomous systems. However, unlike OSPF, BGP4 is not a dynamic protocol. Routing decisions are based on policy. Routes are predetermined and remain relatively stable. BGP4 must be run over a reliable transport protocol. Since TCP is used on most routers and hosts it is used as BGP4s transport protocol. The specification for BGP4 can be found in (Rekhter and Li, 1995). Specifics on implementation of BGP4 in the Internet can be found in (Rekhter and Gross, 1995). Except where specific reference is made to the ADNS implementation of OSPF, this description is a consolidation of relevant sections of these RFCs. The discussion of determining route preferences is consolidated from (Rekhter and Gross, 1995) all other portions are from (Rekhter and Li, 1995).

b. BGP4 Message Types

There are four different message types used by this protocol to communicate between BGP4 hosts.

- Open. This is the first message sent by both ends of a connection. In addition to fields that identify the sending router and its associated AS this message also contains a Hold Time field. Hold Time is the number of seconds allowed between receipt of Update or Keep Alive messages before a link will be considered down.
- Update. This message type is used to transfer the routing table information

between two routers. The message format allows for the transfer of a single feasible route to a destination or to remove unfeasible routes. One message can be constructed to perform both functions.

- Notification. Notifications are sent to indicate an error condition has occurred. The connection along which the message is sent is closed immediately after receipt of the Notification. A Notification will be generated as a result of errors in message content or as a result of the hold timer expiring.
- Keep Alive. A Keep Alive message is used to maintain the open status of a connection. One is sent in response to a valid Open message. When no other messages (i.e., Updates) are being sent a Keep Alive will be generated to maintain the link active. Keep Alive messages are normally sent at about one third of the Hold Time Interval.

c. Operation

The first step in the routing process is the establishment of a TCP connection between the source and destination. Next the entire BGP routing table is sent across the link. Because BGP4 does not require periodic refreshing of the routing table the host must maintain the received table for the duration of the connection. Updates to the table are generated when changes are made.

Once the routing table has been sent the connection is maintained open through the use of periodic Keep Alive messages or Updates. Data is passed via the advertised route to its destination.

d. Routing Decision Process

Each router receives route information from other BGP4 routers via

Updates. This routing information is maintained in a database in the router. The router then applies a set of decision rules to this data to determine its preferred route to a particular destination. The decision process occurs in three phases. The ultimate output of the decision process is a table of routes that are to be advertised to other BGP4 routers.

Phase one involves determining the degree of preference associated with routes received from other BGP4 routers. Upon receipt of an Update message the router will invoke the preference policy implemented in the router. The policy is determined locally for each router and is implemented in the form of configuration information in the router. In general this preference decision can be based on path information or other policy or a combination of both. Path information can include such things as AS count, which is the number of systems that must be traversed to reach the destination. Policy can be used to avoid certain links because of known problems such as reliability or stability. If there are multiple BGP4 routers in an AS they will all invoke the same set of policies. Based on these policies they must internally agree on which router will be advertised to neighboring BGP4 routers as the gateway to that AS.

Phase two evaluates routes to select the preferred route to be advertised to other systems. Once phase one is completed every route to a specific destination is compared and the route with the highest preference is selected. If there is only one route to a particular destination no decision is required and that route is then selected. The result of this phase is a table of containing one preferred route to each reachable destination.

Phase three involves the passing of the results of this process to other BGP4 routers. This is accomplished through the use of Update messages.

LIST OF REFERENCES

Casey, R., Naval Command Control and Ocean Surveillance Center (NRaD Code D8205), *ADNS Implementation Working Paper*, NRaD, July 15, 1997.

Casey, R., Naval Command Control and Ocean Surveillance Center (NRaD Code D8205) and M. Stell, Naval Command Control and Ocean Surveillance Center (NRaD Code D824), *Autonomous System Implementation for Navy Afloat Forces*, June 26, 1997.

Chief of Naval Operations (OPNAV), "Copernicus...Forward", OPNAV Space-Command and Control-Information Warfare Strategic Planning Office (N6C), June 1995

Department of the Navy, "Forward...From the Sea", Navy Public Affairs Library, http://www.chinfo.navy.mil/navpalib/policy/fromsea/forward.txt, November 9, 1994

Johnson, K., SAIC, *Automated Digital Network System LBG Training Material (Draft)*, Naval Command Control and Ocean Surveillance Center (NCCOSC RDT&E Div.), June 30, 1997.

JMCOMS (SPAWAR PMW-176) Brief, Copernicus Requirements Working Group, http://c4iweb.nosc.mil/crwgdoc/, May 1997

Joint Chiefs of Staff, "C4I for the Warrior", CJCS, 12 June 1993.

Moy, J., "Multicast Extension to OSPF", Internet Request for Comments, http//ds.internic.net/rfc/rfc1584.txt1584, March, 1994.

Moy, J., "OSPF Version 2", Internet Request for Comments 2178, http//ds.internic.net/rfc/rfc2178.txt, July, 1997.

Naval Command Control and Ocean Surveillance Center (NRaD Code 80*) Automated Integrated Communications System (AIC) Network Management Architecture for the Advanced Digital Network System (ADNS)*, NRaD, September 30, 1996.

PIAC Brief, Copernicus Requirements Working Group, http://c4iweb.nosc.mil/crwgdoc/

Rekhter, Y. and P. Gross, "Application of the Border Gateway Protocol in the Internet", Internet Request for Comments 1772, http//ds.internic.net/rfc/rfc1772.txt, March, 1995.

Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", Internet Request for Comments 1771, http//ds.internic.net/rfc/rfc1771.txt, March, 1995.

Space and Naval Warfare Systems Command (SPAWAR), PMW-176, *JMCOMS Master Plan*, Version 1.0,SPAWAR PMW-176, March 17, 1997.

Space and Naval Warfare Systems Command (SPAWAR), PMW-176, *Joint Maritime Communications System (JMCOMS) Integrated Network Management (INM) Technical Approach*, draft revision 2, April 21, 1997.

INITIAL DISTRIBUTION LIST

No. of Copies

1. Defense Technical Information Center 2
   8725 John J. Klingman Rd., STE 0944
   Ft. Belvoir, VA 22060-6218

2. Dudley Knox Library 2
   Naval Postgraduate School
   411 Dyer Rd.
   Monterey, CA 93943-5101

3. Joe Macker 1
   Code 5544
   Center for High Assurance Computer Systems
   Navy Research Lab
   Washington, D.C. 20375

4. Roger Casey 1
   NRaD (Code D8205)
   271 Santa Catalina Blvd.
   San Diego, CA 92152

5. Brain Clingerman 1
   SPAWAR, PMW-176
   53560 Hull St.
   San Diego, CA 92152

6. Mike Sovereign 1
   HQ CINCPAC, J56
   Box 64015
   Camp H.M. Smith, HI 96861-4105

7. Professor Rex Buddenberg 1
   Naval Postgraduate School
   Code SM/BU

8. LCDR James A. Sullivan 1
   7 Burr Road
   East Northport, NY 11731

66 553NPS TH 3734
11/99 22527-106 MILES

DUDLEY KNOX LIBRARY
3 2768 00366398 0